Zero Data Loss Recovery Appliance
Owner's Guide

Release 19.2
F16639-06
September 2021

ORACLE

Zero Data Loss Recovery Appliance Owner's Guide, Release 19.2

F16639-06

Contributors: Glenn Maxey

Primary Author: Terence Buencamino

Contributors: Donna Carver, Doug Archambault, Andrew Babb, Tim Chien, Sam Corso, Donna Cooksey, Caroline Johnston, Leslie Keller, Lawrence To, Ericka Washington

Contents

8 Configuring a Recovery Appliance Rack

9 Installing the Recovery Appliance Software

13 Reconfiguring the Hardware

Part IV Appendixes

A RACLI Overview

B Error Messages

C Replacement Units

D In-Rack Cabling Tables for Recovery Appliance X9M

Index

List of Figures

List of Tables

Preface

Welcome to *Zero Data Loss Recovery Appliance Owner's Guide*.

This preface contains the following topics:

Audience

This guide is intended for customers and those responsible for data center site planning, configuration, and maintenance of Zero Data Loss Recovery Appliance, commonly known as Recovery Appliance.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.
$ prompt	The dollar sign ($) prompt indicates a command run as the `oracle` user.
# prompt	The pound (#) prompt indicates a command that is run as the `root` user.

Changes in Oracle Zero Data Loss Recovery Appliance Release 19.2.1.1.2

The following are changes in *Oracle Zero Data Loss Recovery Appliance* for ZDLRA release 19.2.1.1.2.

- Updates and additions to support the X9M Compute Server, including cabling and descriptions.
- The security information from Exadata was included in a new chapter and updated to reflect the requirements for the Recovery Appliance. User Security on Recovery Appliance. The default user accounts for the Recovery Appliancewere moved from an appendix into this chapter to match Exadata.
- The Oracle Exadata Deployment Assistant (OEDA) was included in a new chapter and updated to reflect the requirements for the Recovery Appliance. Using the Deployment Assistant .
- Updated commands for RACLI.
  - racli add host
  - racli alter cloud_user
  - racli configure archive_group
  - racli disable aide
  - racli enable aide
  - racli list cloud_location
  - racli list host
  - racli list grid_home
  - racli list oracle_home
  - racli remove grid_home
  - racli remove host
  - racli remove oracle_home
  - racli remove osb_pieces
  - racli remove oracle_home
  - racli status aide
  - racli update aide

Part I
Preinstallation

Part I provides an introduction to Zero Data Loss Recovery Appliance. It describes the preinstallation procedures that you must complete before receiving shipment, to prevent delays with the installation.

Part I contains the following chapters:

- Introduction to Recovery Appliance
- #unique_19
- Preparing Your Networks for Recovery Appliance
- Setting Up Auto Service Request
- Using the Deployment Assistant

Site Requirements for Recovery Appliance and Oracle Exadata Storage Expansion Rack

This chapter describes the site requirements for Recovery Appliance and Oracle Exadata Storage Expansion Rack.

> **Note:**
>
> For ease of reading, the name "ZDLRA Rack" is used when information refers to both Recovery Appliance and Oracle Exadata Storage Expansion Rack.

Related Topics

- #unique_32

Ensuring That the Site is Ready

Before ZDLRA Rack is delivered to the site, the following tasks must be performed to ensure the site is ready:

Task 1 Review Site Requirements
Review the site requirements in this chapter to understand the requirements for ZDLRA Rack.

Task 2 Understand the Networking Requirements and Options
Review #unique_34 to discern the networking requirements and configuration decisions that you must make.

Task 3 Run Oracle Exadata Deployment Assistant (OEDA)
Run OEDA to create the necessary configuration files.
See Using the Deployment Assistant .

Task 4 Configure the Network
Configure the network using the files from OEDA. This includes registering the networks in the Domain Name System (DNS), assigning IP addresses, and configuring data center switches and firewalls.

Task 5 Prepare the Site Based on Requirements
Prepare the site based on the requirements, such as, install the network cables and power supplies, prior to the arrival of ZDLRA Rack.
Review the procedures in Installing the Recovery Appliance at the Site

General Environmental Requirements

The following sections describe the general environmental requirements for ZDLRA Racks.

General Environmental Requirements for ZDLRA Rack X6 and Later

Starting with X6 models of Recovery Appliance and Oracle Exadata Storage Expansion Rack, use Oracle Exadata Configuration Assistant (OECA) to determine environmental requirements, such as size, weight, acoustic level, power, cooling, and airflow. This includes all X6, X7, X8, X8M, X9M, and later models.

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

Space Requirements

For all Recovery Appliance X7, X8, X8M, X9M, and later models, use Oracle Exadata Configuration Assistant (OECA) to determine the size of each rack.

In addition to the rack height, 914 mm (36 inches) of space is required above the rack height for maintenance access, as described in Table 21.

The space surrounding the cabinet must not restrict the movement of cool air between the air conditioner and the front of the systems within the cabinet, or the movement of hot air coming out of the rear of the cabinet.

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

Space Requirements for Racks up to Recovery Appliance X6

All racks up to Recovery Appliance X6-2 use the same hardware rack, and have the same space requirements.

The space requirements are as follows:

- Height: 1998 mm (78.66 inches)
- Width: 600 mm with side panels (23.62 inches)

- Depth: 1200 mm (47.24 inches)

The minimum ceiling height for the cabinet is 2912 mm (114.65 inches), measured from the true floor or raised floor, whichever is higher. This includes an additional 914 mm (36 inches) of space required above the rack height for maintenance access, as described in Table 21. The space surrounding the cabinet must not restrict the movement of cool air between the air conditioner and the front of the systems within the cabinet, or the movement of hot air coming out of the rear of the cabinet.

Receiving, Unpacking, and Access Route Requirements

Before your ZDLRA Rack arrives, ensure that the receiving area is large enough for the package.

Use the following package dimensions for ZDLRA Rack:

- Shipping height: 2159 mm (85 inches)
- Shipping width: 1219 mm (48 inches)
- Shipping depth: 1575 mm (62 inches)

If your loading dock meets the height and ramp requirements for a standard freight carrier truck, then you can use a pallet jack to unload the rack. If the loading dock does not meet the requirements, then you must provide a standard forklift or other means to unload the rack. You can also request that the rack be shipped in a truck with a lift gate.

When ZDLRA Rack arrives, leave the rack in its shipping packaging until it arrives at its installation site. Use a conditioned space to remove the packaging material to reduce particles before entering the data center. The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.

Allow enough space for unpacking it from its shipping cartons. Ensure that there is enough clearance and clear pathways for moving ZDLRA Rack from the unpacking location to the installation location.

> **Caution:**
>
> Prior to moving the rack, always make sure that all four leveling and stabilizing feet are raised and out of the way.

Table 19 Access Route Requirements

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	2184 mm (86 inches)	2040 mm (80.32 inches)
Minimum door width	1270 (50 inches)	640 mm (25.19 inches)
Minimum elevator depth	1625.6 mm (64 inches)	1240 mm (48.82 inches)
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity	1134 kg (2500 lbs)	1134 kg (2500 lbs)

Rack Weights for Recovery Appliance

Review the weight of your ZDLRA Rack to ensure that it can be delivered and located safely.

For all Recovery Appliance X6, X7, X8, X8M, X9M, and later models, use Oracle Exadata Configuration Assistant (OECA) to determine the net rack weight. In addition to the net rack weight, allow approximately 180 kg (400 lbs) for packaging materials.

For earlier models, use the following table to determine the rack weight:

Table 20 Shipping Weight

Recovery Appliance	Full Rack	Minimum Configuration
X7 Weight	899.9 kg (1983.9 lbs)	385.7 kg (850.4 lbs)
X6 and X5 Weight	875.7 kg (1930.5 lbs)	381.2 kg (840.5 lbs)
X4 Weight	866 kg (1909 lb)	469 kg (1034 lb)

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

Maintenance Access Requirements

The maintenance area must be large enough for ZDLRA Rack, and have the required access space.

The required space to remove the side panels is 675.64 mm (26.6 inches). Side panel access is not required for normal maintenance.

Front access space requirement is typically the length of the longest chassis supported in the rack, plus the distance the rails need to travel out to install and remove equipment from the front.

Rear access space requirement is the space needed for a person working behind the ZDLRA Rack.

Top access space requirement is for a person standing on a ladder to be able service components installed in the top of the rack.

Table 21 Maintenance Access Requirements for ZDLRA Rack (All Models)

Location	Maintenance Access Requirement
Rear maintenance	914 mm (36 inches)
Front maintenance	1232 (48.5 inches)
Top maintenance	914 mm (36 inches)

> **Note:**
>
> - In raised floor data centers, an adjacent tile at the rear of the rack must be able to be opened to access any cabling going under the floor.
> - ZDLRA Racks can be placed side-by-side when interconnecting the racks.

Flooring Requirements

ZDLRA Rack may be installed on raised floor or solid floor environments.

The site floor or the raised flooring must be able to support the total weight of ZDLRA Rack.

Table 22 Floor Load Requirements for ZDLRA Rack

Description	Requirement
Maximum allowable weight of installed rack equipment	952.5 kg (2100 lbs)
Maximum allowable weight of installed power distribution units	52.16 kg (115 lbs)
Maximum dynamic load (maximum allowable weight of installed equipment including PDUs)	1004.66 kg (2215 lbs)

> **Note:**
>
> For raised floors, a floor tile with a cutout at the rear of the Exadata rack is needed for routing power and network cables into the rack from under the raised floor. The cutout is located between the two rear leveling feet. See Figure 5-3 and Figure 5-5.

Electrical Power Requirements

The ZDLRA Rack can operate effectively over a wide range of voltages and frequencies. However, each rack must have a reliable power source.

Damage may occur if the ranges are exceeded. Electrical disturbances such as the following may damage ZDLRA Rack:

- Fluctuations caused by brownouts
- Wide and rapid variations in input voltage levels or in input power frequency
- Electrical storms
- Faults in the distribution system, such as defective wiring

To protect ZDLRA Rack from such disturbances, you should have a dedicated power distribution system, power-conditioning equipment, as well as lightning arresters or power cables to protect from electrical storms.

PDU Power Requirements

Each rack has two pre-installed power distribution units (PDUs). The PDUs accept different power sources. You must specify the type of PDU that is correct for your ZDLRA Rack and data center.

The following are the PDUs that Oracle Zero Data Loss Recovery Appliance currently supports or has supported in the past:

Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Table 23 Specifications for Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Specification	Required values
Voltage	200-240 VAC
Frequency	50/60 Hz
Current	24 A maximum per input
Power rating	15 kVA
Output current	72 A (3 x 24 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A

Materials Required

The following are needed to connect an Oracle Exadata rack to a low-voltage three-phase power source:

- 6 power cords for two PDUs, each rated for 30 amperes at 200-240 VAC
- 6 receptacles for 2 PDUs, each must be 15kVA, with three 30A/250V 2-pole/3-wire NEMA L6-30P plugs

Figure 1 Low-voltage Single Phase Power Connector for North America, South America, Japan and Taiwan

Low-voltage 15 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Table 24 Specifications for Low-voltage 15 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Specification	Required values, per PDU
Voltage	190-220 VAC
Frequency	50/60 Hz
Current	40 A maximum per phase
Power rating	14.4 kVA
Output current	69.3 A (3 x 23.1A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	IEC309-3P4W-IP67 60A 250VAC 3ph (Hubbell equivalent is HBL460R9W)

Materials Required

The following are needed to connect an Oracle Exadata system to a low-voltage three-phase power source:

- 2 power cords for two PDUs, 60 amperes at 190-220 VAC three-phase
- 2 receptacles to connect the PDUs to 2 IEC 60309 60A 4-pin 250VAC 3ph IP67 data center receptacles

Figure 2 Low-voltage Three Phase Power Connector for North America, South America, Japan and Taiwan

High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 25 Specifications for High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Required values, per PDU
Voltage	220-240 VAC

Table 25 (Cont.) Specifications for High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Required values, per PDU
Frequency	50/60 Hz
Current	24 A maximum per input
Power rating	15 kVA
Output current	72 A (3 x 24 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	15kVA, IEC309-2P3W-IP44 32A 250VAC (Hubbell equivalent is HBL332R6W

Note:

The high-voltage 15kVA single phase PDU is no longer available. The high-voltage 22kVA is a compatible replacement with the same receptacle and breaker requirements

Figure 3 High-voltage Single Phase Power Connector

High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 26 Specifications for High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Required values, per PDU
Voltage	220/380-240/415 VAC 3-phase
Frequency	50/60 Hz
Current	25 A maximum per phase
Power rating	14.4 kVA

Table 26 (Cont.) Specifications for High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Required values, per PDU
Output current	62.7 A (3 x 20.9 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 1-pole circuit breaker)	20A
Data center receptacle	15kVA, three phase, 5-pin, IEC 60309 32A, 5-pin 230/400V, three phase IP44

Materials Required

The following are needed to connect an Oracle Exadata system to a high-voltage three-phase power source:

- 2 power cords for two PDUs, 25 amperes at 220/380-240/415 VAC three-phase
- 2 receptacles to connect the PDUs to 2 IEC 60309 32A 5-pin 230/400VAC 3ph IP44 data center receptacles

Figure 4 High-voltage Three Phase Power Connector

Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Table 27 Specifications for Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Specifications	Required values, per PDU
Voltage	200-240 VAC
Frequency	50/60 Hz
Current	36.8 A maximum per input / 110.4 A Maximum per PDU
Power rating	22 kVA
Output current	110.4A (3x 36.8)

Table 27 (Cont.) Specifications for Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Specifications	Required values, per PDU
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	Hubbell CS8269 or CS8264

Materials Required

The following are needed to connect an Oracle Exadata system to a low-voltage single-phase power source:

- 3 power cords for two PDUs
- 6 receptacles to connect the PDUs

The following image shows the low-voltage, 22 kVA single phase PDU power connector for North America, South America, Japan and Taiwan. This power connector is different from the power connector for the low-voltage 15 kVA single phase PDU.

Figure 5 Low-voltage 22 kVA Single Phase Power Connector

High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 28 Specifications for High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specifications	Required values, per PDU
Voltage	220-240 VAC
Frequency	50/60 Hz
Current	32 A maximum per input
Power rating	22 kVA
Output current	96 A (3 x 32 A)
Outlets	42 x C13; 6 x C19

Table 28 (Cont.) Specifications for High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specifications	Required values, per PDU
Outlet groups	6
Group protection (UL489 1-pole circuit breaker)	20A
Data center receptacle	IEC309-2P3W-IP44 32A 250VAC (Hubbell equivalent is HBL332R6W)

Materials Required

> **Note:**
>
> The high-voltage 15kVA single phase PDU is no longer available. The high-voltage 22kVA is a compatible replacement with the same receptacle and breaker requirements

The following are needed to connect an Oracle Exadata system to a high-voltage one-phase power source:

- 6 power cords for two PDUs, each rated for 25 amperes at 220/380-240/415 VAC 1-phase voltage
- 6 receptacles to connect the PDUs

Figure 6 High-voltage Single Phase Power Connector

Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Table 29 Specifications for Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Specifications	Required values, per PDU
Voltage	200-208 VAC 3-phase
Frequency	50/60 Hz

Table 29 (Cont.) Specifications for Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Specifications	Required values, per PDU
Current	34.6 A maximum per phase
Power rating	25 kVA
Output current	120 A (6 x 20 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	IEC309-3P4W-IP67 60A 250VAC 3ph (Hubbell equivalent is HBL460R9W)

Materials Required

The following are needed to connect an Oracle Exadata system to a low-voltage three-phase power source:

- 4 power cords for two PDUs
- 4 receptacles to connect the PDUs

High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 30 Specifications for High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specifications	Required values, per PDU
Voltage	220/380-240/415 VAC 3-phase
Frequency	50/60 Hz
Current	18.1 A maximum per phase
Power rating	25 kVA
Output current	108.6A (6 x 18.1 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 1-pole circuit breaker)	20A
Data center receptacle	IEC309-4P5W-IP44 32A 400VAC 3ph (Hubbell equivalent is HBL532R6W)

Materials Required

The following are needed to connect an Oracle Exadata system to a high-voltage three-phase power source:

- 4 power cords for two PDUs
- 4 receptacles to connect the PDUs

Facility Power Requirements

To prevent catastrophic failures, design the input power sources to ensure adequate power is provided to the PDUs.

Use dedicated AC breaker panels for all power circuits that supply power to the PDU. When planning for power distribution requirements, balance the power load between available AC supply branch circuits. In the United States of America and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.

> **Note:**
>
> Electrical work and installations must comply with applicable local, state, or national electrical codes. Contact your facilities manager or qualified electrician to determine what type of power is supplied to the building.

PDU power cords are 4 meters (13.12 feet) long, and 1 to 1.5 meters (3.3 to 4.9 feet) of the cord is routed within the rack cabinet. The installation site AC power receptacle must be within 2 meters (6.6 feet) of the rack.

Circuit Breaker Requirements

If computer equipment is subjected to repeated power interruptions and fluctuations, then it is susceptible to a higher rate of component failure.

You are responsible for supplying the circuit breakers. One circuit breaker is required for each power cord. In addition to circuit breakers, provide a stable power source, such as an uninterruptible power supply (UPS) to reduce the possibility of component failures.

Use dedicated AC breaker panels for all power circuits that supply power to the server. Servers require electrical circuits be grounded to the Earth.

> **Note:**
>
> Electrical work and installations must comply with applicable local, state, or national electrical codes.

Electrical Grounding Guidelines

The cabinets for ZDLRA Rack are shipped with grounding-type power cords.

- Always connect the cords to grounded power outlets.
- Check the grounding type, because different grounding methods are used depending on your location.
- Refer to documentation such as IEC documents for the correct grounding method.

- Ensure that the facility administrator or qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.

Temperature and Humidity Requirements

Excessive internal temperatures may result in full or partial shut down of Recovery Appliance.

Airflow through ZDLRA Rack is from front to back. Refer to General Environmental Requirements for information on cooling and airflow.

> **Note:**
>
> Studies have shown that temperature increases of 10 degrees Celsius (15 degrees Fahrenheit) above 20 degrees Celsius (70 degrees Fahrenheit) reduce long-term electronics reliability by 50 percent.

The following table lists the temperature, humidity and altitude requirements for operating and non-operating machines.

Table 31 Temperature, Humidity and Altitude Requirements

Condition	Operating Requirement	Non-operating Requirement	Optimum
Temperature	5 to 32 degrees Celsius (41 to 89.6 degrees Fahrenheit)	-40 to 70 degrees Celsius (-40 to 158 degrees Fahrenheit).	For optimal rack cooling, data center temperatures from 21 to 23 degrees Celsius (70 to 74 degrees Fahrenheit)
Relative humidity	10 to 90 percent relative humidity, non-condensing	Up to 93 percent relative humidity.	For optimal data center rack cooling, 45 to 50 percent, non-condensing
Altitude	3048 meters (10000 feet) maximum	12000 meters (40000 feet).	Ambient temperature is reduced by 1 degree Celsius per 300 m above 900 m altitude above sea level

Set conditions to the optimal temperature and humidity ranges to minimize the chance of downtime due to component failure. Operating ZDLRA Rack for extended periods at or near the operating range limits, or installing it in an environment where it remains at or near non-operating range limits could significantly increase hardware component failure.

The ambient temperature range of 21 to 23 degrees Celsius (70 to 74 degrees Fahrenheit) is optimal for server reliability and operator comfort. Most computer equipment can operate in a wide temperature range, but near 22 degrees Celsius (72 degrees Fahrenheit) is desirable because it is easier to maintain safe humidity levels. Operating in this temperature range provides a safety buffer in the event that the air conditioning system goes down for a period of time.

The ambient relative humidity range of 45 to 50 percent is suitable for safe data processing operations. Most computer equipment can operate in a wide range (20 to 80 percent), but the range of 45 to 50 percent is recommended for the following reasons:

- Optimal range helps protect computer systems from corrosion problems associated with high humidity levels.
- Optimal range provides the greatest operating time buffer in the event of air conditioner control failure.
- This range helps avoid failures or temporary malfunctions caused by intermittent interference from static discharges that may occur when relative humidity is too low.

Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent. It is not difficult to maintain humidity in a data center because of the high-efficiency vapor barrier and low rate of air changes normally present.

Ventilation and Cooling Requirements

Always provide adequate space in front and behind the rack to allow for proper ventilation.

Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Rack-mountable servers and equipment typically draw cool air in through the front of the rack and let out warm air out the rear of the rack. There is no air flow requirement for the left and right sides due to front-to-back cooling.

If the rack is not completely filled with components, then cover the empty sections with filler panels. Gaps between components can adversely affect air flow and cooling within the rack.

Relative humidity is the percentage of the total water vapor that can exist in the air without condensing, and is inversely proportional to air temperature. Humidity goes down when the temperature rises, and goes up when the temperature drops. For example, air with a relative humidity of 45 percent at a temperature of 24 degrees Celsius (75 degrees Fahrenheit) has a relative humidity of 65 percent at a temperature of 18 degrees Celsius (64 degrees Fahrenheit). As the temperature drops, the relative humidity rises to more than 65 percent, and water droplets are formed.

Air conditioning facilities usually do not precisely monitor or control temperature and humidity throughout an entire computer room. Generally, monitoring is done at individual points corresponding to multiple exhaust vents in the main unit, and other units in the room. Special consideration should be paid to humidity when using underfloor ventilation. When underfloor ventilation is used, monitoring is done at each point close to an exhaust vent. Distribution of the temperature and humidity across the entire room is uneven.

Oracle Exadata Racks have been designed to function while mounted in a natural convection air flow. The following requirements must be followed to meet the environmental specification:

- Ensure there is adequate air flow through the server.
- Ensure the server has front-to-back cooling. The air inlet is at the front of the server, and the air is let out the rear.
- Allow a minimum clearance of 1219.2 mm (48 inches) at the front of the server, and 914 mm (36 inches) at the rear of the server for ventilation.

Use perforated tiles, approximately 400 CFM/tile, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, as long as cold air from the tiles can flow into the rack. Inadequate cold air flow could result in a higher inlet temperature in the

servers due to exhaust air recirculation. The following table describes the recommended number of floor tiles:

Type of Oracle Exadata Database Machine or Oracle Exadata Storage Expansion Rack	Recommended Number of Floor Tiles
Oracle Exadata full rack systems	4
Oracle Exadata half rack systems	3
Oracle Exadata quarter and eighth rack systems	1

Figure 7 shows a typical installation of the floor tiles for a full rack.

Figure 7 Typical Data Center Configuration for Perforated Floor Tiles

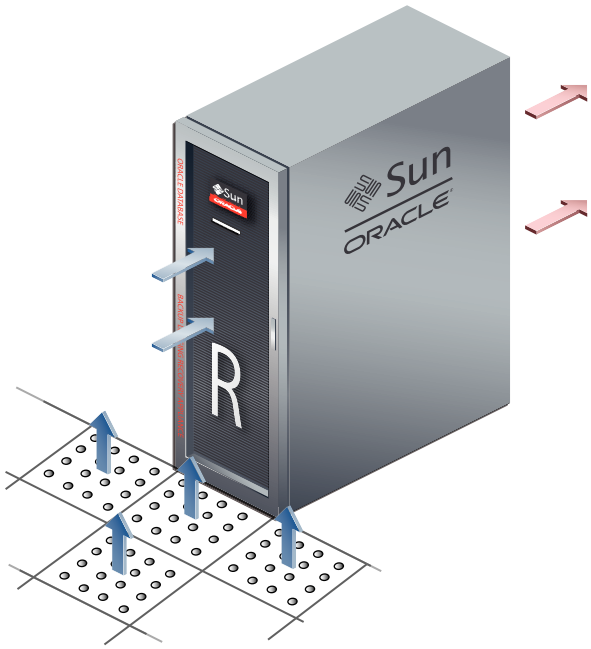

Network Connection and IP Address Requirements for ZDLRA Rack

This section describes the requirements to connect the Oracle Zero Data Loss Recovery Appliance to your existing network infrastructure.

Network Connection Requirements for ZDLRA Rack

Prior to installation, network cables must be run from your existing network infrastructure to the installation site. The requirements to connect Recovery Appliance to your existing network infrastructure are as follows:

Network	Connection Requirements
Administration Network	• Mandatory: 1 uplink for the Cisco Management Switch. • Recommended: 2 uplinks for remote monitoring of the power distribution units (PDUs). • Optional: 1 uplink for ILOM on each database server or storage server if using a separate ILOM network.
Client Network	• For bonded network configurations: 2 uplinks for each database server. • For non-bonded network configurations: 1 uplink for each database server. **Note:** Non-bonded network configurations are not supported on Recovery Appliance X7 and later systems.
Additional Network(s)	• For bonded network configurations: 2 uplinks for each database server and additional network. • For non-bonded network configurations: 1 uplink for each database server and additional network. **Note:** Non-bonded network configurations are not supported on Recovery Appliance X7 and later systems.
RDMA Network Fabric (Private Network)	No uplinks are required. The RDMA Network Fabric is completely self-contained within Recovery Appliance.

Related Topics

- #unique_34

DNS Configuration for ZDLRA Rack

Prior to receiving your ZDLRA Rack, you must use Oracle Exadata Deployment Assistant (OEDA) to generate files that drive the configuration of the system. The system configuration files contain network host names and IP addresses for the management network, client network, and additional public networks, which should be registered in Domain Name System (DNS) prior to initial configuration. In particular, all public addresses, single client access name (SCAN) addresses, and VIP addresses should be registered in DNS prior to installation.

The assistant-generated file defines the SCAN as a single name with three IP addresses on the client access network. The three SCAN addresses provide service access for clients to Oracle Zero Data Loss Recovery Appliance. Configure DNS for round robin resolution for the SCAN name to these three SCAN addresses.

All addresses registered in DNS must be configured for both forward resolution and reverse resolution. Reverse resolution must be forward confirmed (forward-confirmed reverse DNS) such that both the forward and reverse DNS entries match each other.

IP Address Requirements for ZDLRA Rack

The Recovery Appliance requires a dedicated allocation of host names and IP addresses. The number of required IP addresses primarily depends on the size of the system.

The detailed network configuration, including host names and IP addresses, is generated from information supplied to Oracle Exadata Deployment Assistant (OEDA).

You should configure the IP addresses in your existing networks after you run OEDA. All IP addresses in the configuration file must be unassigned at the time of initial configuration. In addition, all IP addresses must be statically-assigned IP addresses, not dynamically-assigned (DHCP) addresses.

All RDMA Network Fabric IP addresses must be in the same subnet, with a minimum subnet mask of 255.255.240.0 (or /20). The subnet mask chosen should be wide enough to accommodate possible future expansion of the Recovery Appliance and the internal network.

You can use the information in the following sections to calculate the IP address requirement for Recovery Appliance. However, ensure that you use the configuration information generated in OEDA as your authoritative guide.

Rack-Level IP Address Requirements

The following table outlines the rack-level IP address requirements for Recovery Appliance. These addresses are required regardless of the configuration options that are selected for the system.

Network	IP Address Requirements
Administration Network	• 1 IP address for the administration network interface on each compute server or storage server. • 1 IP address for the ILOM network interface on each database server or storage server. • 1 IP address for the management interface on each RDMA Network Fabric switch. • 1 IP address for the management interface on the Cisco Management Switch switch. • 1 IP address for the management interface in each power distribution unit (PDU).
RDMA Network Fabric (Private Network)	2 IP addresses for each database server or storage server.

Related Topics

- Single Client Access Name (SCAN)
- #unique_65

1

Introduction to Recovery Appliance

This chapter describes the features and hardware components of Oracle's Zero Data Loss Recovery Appliance (Recovery Appliance).

This chapter contains these sections:

About Recovery Appliance

Recovery Appliance is a backup and recovery solution that is specifically engineered to protect your Oracle Database data. It offers a massively scalable repository that records all delta changes (that is, differences) across hundreds to thousands of databases.

The Database Delta Push technology enables Oracle DBAs to run one full Recovery Manager (RMAN) backup, and thereafter run only incremental backups, from each database to the Recovery Appliance. There are no recurring full backups, tape backups, or RMAN backup deletion, validation, and maintenance commands.

In the event of a disaster, DBAs can recover any database to any point in time in the recovery window or within subseconds of the current time. The Database Delta Store virtual full backup technology constructs physical backups from the resident incremental blocks, in response to RMAN database restore requests.

You can easily and quickly replicate backups on a local Recovery Appliance over a secure network to a remote Recovery Appliance for protection against server or site outage. Only the changed blocks are replicated. Virtual full backups on the remote Recovery Appliance are available for database restore operations.

Autonomous tape archives further protect against site disasters. Recovery Appliance automates the copying of virtual full or incremental backups to tape, for long-term retention and archival purposes. Recovery Appliance is preinstalled with tape backup software and support for an optional 16 Gb fiber channel card on each compute server, to connect directly to the tape hardware.

You can manage this entire infrastructure with end-to-end visibility into the status of data protection to ensure that data can be recovered successfully at any time. The primary administrative tools are Oracle Enterprise Manager Cloud Control and the RMAN command interface.

Recovery Appliance Rack Components

Recovery Appliance consists of compute servers, storage servers, and the network components to connect to your network. Oracle offers a flexible configuration for Recovery Appliance, so that you can acquire the resources that you need, and expand incrementally when you need more. You can configure up to 18 full racks into a single cluster.

Table 1-1 lists the components of a Recovery Appliance X8M rack.

Table 1-1 Components of a Recovery Appliance X8M Rack

Quantity	Component
2	Oracle Server X8M-2 compute servers
3 to 18	Oracle Storage Server X8M-2 storage servers
2	36-Port Cisco Nexus 9336c Ethernet Switches
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

Table 1-2 lists the components of a Recovery Appliance X8-2 rack.

Table 1-2 Components of a Recovery Appliance X8-2 Rack

Quantity	Component
2	Oracle Server X8-2 compute servers
3 to 18	Oracle Storage Server X8-2 storage servers
2	36-port Sun Datacenter InfiniBand Switch
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

#unique_23/unique_23_Connect_42_COMPONENTSOFARECOVERYAPPLIANCEX7RAC-36AC788C lists the components of a Recovery Appliance X7 rack.

Table 1-3 Components of a Recovery Appliance X7 Rack

Quantity	Component
2	Oracle Server X7-2 compute servers
3 to 18	Oracle Server X7-2L storage servers
2	36-port Sun Datacenter InfiniBand Switch

Table 1-3 (Cont.) Components of a Recovery Appliance X7 Rack

Quantity	Component
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

Table 1-4 lists the components of a Recovery Appliance X6 rack.

Table 1-4 Components of a Recovery Appliance X6 Rack

Quantity	Component
2	Oracle Server X6-2 compute servers
3 to 18	Oracle Server X6-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	16 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports

Table 1-5 lists the components of a Recovery Appliance X5 rack.

Table 1-5 Components of a Recovery Appliance X5 Rack

Quantity	Component
2	Oracle Server X5-2 compute servers
3 to 18	Oracle Server X5-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	16 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports

Table 1-6 lists the components of a Recovery Appliance X4 rack.

Table 1-6 Components of a Recovery Appliance X4 Rack

Quantity	Component
2	Sun Server X4-2 compute servers
3 to 14	Sun Server X4-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	16 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports

Compute Server Components

Table 1-7 lists the components of an individual Oracle X8M-2 compute server.

Table 1-7 X8M-2 Compute Server Components

Quantity	Description
2	24-core Xeon 8160 processors (2.4 GHz)
8	384 GB RAM
2	100Gb/s RDMA Network Fabric Card
1	10/100/1000BASE-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
1	1 GB copper Ethernet Port (mgmt)
2	On-board 10 Gigabit copper Base-T Ethernet ports
2	On-board 10/25 Gigabit optical Ethernet ports
2	Optional PCIe card 10/25 Gigabit optical Ethernet ports Sun Dual 10/25 Gb Ethernet SFP+/28 PCIe 2.0 Low Profile Adapter incorporating Intel 82599 10/25 Gb Ethernet controller and supporting pluggable SFP+ Transceivers. ROHS-5
2	32 Gb Fibre Channel Ports Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity (optional)
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 2 GB cache

Table 1-8 lists the components of an individual Oracle X8-2 compute server.

Table 1-8 Oracle X8-2 Compute Server Components

Quantity	Description
2	24-core Xeon 8160 processors (2.4 GHz)
8	384 GB RAM
2	QDR InfiniBand (40 Gb/s) ports
1	10/100/1000BASE-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
1	1 GB copper Ethernet Port (mgmt)
2	On-board 10 Gigabit copper Base-T Ethernet ports
2	On-board 10/25 Gigabit optical Ethernet ports
2	Optional PCIe card 10/25 Gigabit optical Ethernet ports Sun Dual 10/25 Gb Ethernet SFP+/28 PCIe 2.0 Low Profile Adapter incorporating Intel 82599 10/25 Gb Ethernet controller and supporting pluggable SFP+ Transceivers. ROHS-5

Table 1-8 (Cont.) Oracle X8-2 Compute Server Components

Quantity	Description
2	32 Gb Fibre Channel Ports Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity (optional)
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 2 GB cache

Table 1-9 lists the components of an individual Oracle X7-2 compute server.

Table 1-9 Oracle X7-2 Compute Server Components

Quantity	Description
2	24-core Xeon 8160 processors (2.1 GHz)
8	384 GB RAM
2	QDR InfiniBand (40 Gb/s) ports
1	Oracle Integrated Lights Out Manager (ILOM) for remote management. Ethernet Port
1	1 GB copper Ethernet Port (mgmt)
2	On-board 10 Gigabit copper Base-T Ethernet ports
2	On-board 10/25 Gigabit optical Ethernet ports
2	Optional PCIe card 10/25 Gigabit optical Ethernet ports Sun Dual 10/25 Gb Ethernet SFP+/28 PCIe 2.0 Low Profile Adapter incorporating Intel 82599 10/25 Gb Ethernet controller and supporting pluggable SFP+ Transceivers. ROHS-5
2	32 Gb Fibre Channel Ports Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity (optional)
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 1 GB Supercap-backed Write Cache

Table 1-10 lists the components of an Oracle X6-2 compute server.

Table 1-10 Oracle Server X6-2 Components

Quantity	Description
2	22-core Intel Xeon E5-2699 v4 processors (2.2 GHz)
8	32 GB RAM (256 GB total)
4	600 GB 10K RPM SAS disks
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	1 GbE/10GbE Base-T Ethernet ports

Table 1-10 (Cont.) Oracle Server X6-2 Components

Quantity	Description
2	10 GbE Ethernet SFP+ ports (1 dual-port 10GbE PCIe 2.0 network card based on the Intel 82599 10 GbE controller technology)
1	Ethernet port for Oracle Integrated Lights Out Manager (ILOM) for remote management
1	16 Gb dual-port QLogic PCIe 3.0 HBA card (optional)

Table 1-11 lists the components of an Oracle X5-2 compute server.

Table 1-11 Oracle Server X5-2 Components

Quantity	Description
2	18-core Intel Xeon E5-2699 v3 processors (2.3 GHz)
8	32 GB RAM (256 GB total)
4	600 GB 10K RPM SAS disks
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	1 GbE/10GbE Base-T Ethernet ports
2	10 GbE Ethernet SFP+ ports (1 dual-port 10GbE PCIe 2.0 network card based on the Intel 82599 10 GbE controller technology)
1	Ethernet port for Oracle Integrated Lights Out Manager (ILOM) for remote management
1	16 Gb dual-port QLogic PCIe 3.0 HBA card (optional)

Table 1-12 lists the components of a Sun X4-2 compute server.

Table 1-12 Sun Server X4-2 Components

Quantity	Description
2	12-Core Intel Xeon E5-2697 v2 processors (2.7 GHz)
16	16 GB RAM (256 total)
4	600 GB 10K RPM SAS disks
1	Disk controller HBA with 512 MB battery-backed write cache, and swappable battery backup unit (BBU)
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	1 GbE/10GbE Base-T Ethernet ports
2	10 GbE Ethernet SFP+ ports (1 dual-port 10 GbE PCIe 2.0 network card based on the Intel 82599 10 GbE controller technology)
1	Ethernet port for Oracle Integrated Lights Out Manager (ILOM) for remote management
1	16 Gb dual-port QLogic PCIe 3.0 HBA card (optional)

Storage Server Components

Table 1-13 lists the components of an Oracle X8-2X8-2 and X8M-2 storage server.

Table 1-13 Oracle X8-2 and X8M-2 Server Storage Components

Quantity	Description
2	16-core Intel Xeon 5218 Processors (2.3GHz)
8	192 GB RAM
12	14 TB High Capacity SAS disks
4	6.4 TB flash accelerator PCIe cards
1	Disk controller HBA with 2 GB cache
2	**X8M-2 only:** 100Gb/s RDMA Network Fabric Card, 1 dual-port (PCIe 3.0), both ports active
2	**X8-2 only:** InfiniBand 4X QDR (40 Gb/s) InfiniBand ports, 1 dual-port PCIe 3.0 Host Channel Adapter (HCA), all ports active
2	Embedded Gigabit Ethernet ports
1	10/100/1000BASE-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans

Table 1-14 lists the components of an Oracle X7–2L storage server.

Table 1-14 Oracle Server X7–2L Storage Components

Quantity	Description
2	10-core Xeon 4114 processors (2.0 GHz)
8	192 GB RAM
12	10 TB 7.2 K RPM High Capacity SAS disks
2	3.2 TB non-volatile memory express (NVMe) PCIe 3.0 flash accelerator cards
1	Disk controller HBA with 1 GB supercap-backed write cache
2	36 port QDR InfiniBand (40 Gb/s) ports
2	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans

Table 1-15 lists the components of an Oracle X6-2L storage server.

Table 1-15 Oracle X6-2L Storage Server Components

Quantity	Description
2	10-core Intel Xeon CPU E5-2630 v4 (2.2 GHz)
8	16 GB RAM (128 GB total RAM)
12	8 TB 7.2 K RPM High Capacity SAS disks
4	3.2 TB non-volatile memory express (NVMe) PCIe 3.0 flash accelerator cards
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (PCIe 3.0), both ports active
4	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management

Table 1-16 lists the components of an Oracle Server X5-2L storage server.

Table 1-16 Oracle Server X5-2L Components

Quantity	Description
2	8-core Intel Xeon CPU E5-2630 v3 (2.4 GHz)
4 4	8 GB RAM 16 GB RAM (96 GB total RAM)
12	4 TB 7.2 K RPM High Capacity SAS disks
4	1.6 TB non-volatile memory express (NVMe) PCIe 3.0 flash accelerator cards
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management

Table 1-17 lists the components of a Sun Server X4-2L storage server.

Table 1-17 Sun Server X4-2L Components

Quantity	Description
2	Six-core Intel Xeon E5-2630 v2 processors (2.6 GHz)
4 4	8 GB RAM 16 GB RAM (96 GB total RAM)
12	4 TB 7.2K RPM High Capacity SAS disks
4	800 GB Sun Flash Accelerator F80 PCIe Cards
1	Disk controller HBA with 512 MB battery-backed write cache and swappable BBU

Table 1-17 (Cont.) Sun Server X4-2L Components

Quantity	Description
2	InfiniBand 4 X QDR (40 Gb/s) InfiniBand ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management

Spares Kit Components for Recovery Appliance

The spares kit for Recovery Appliance X7 has the following components:

- 1 x 10 TB 7.2K RPM high capacity SAS disk
- 1 x flash accelerator F320 PCIe card
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

The spares kit for Recovery Appliance X6 has the following components:

- 1 x 8 TB 7.2K RPM high capacity SAS disk
- 1 x flash accelerator F320 PCIe card for high capacity storage servers only
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

The spares kit for Recovery Appliance X5 has the following components:

- 1 x 4 TB 7.2K RPM high capacity SAS disk
- 1 x flash accelerator F160 PCIe card for high capacity storage servers only
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

The spares kit for Recovery Appliance X4 has the following components:

- 1 x 4 TB 7.2K RPM high capacity SAS disk
- 1 x 800 GB Exadata Smart Flash Cache card
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

These spare cables are tied inside the Recovery Appliance rack:

- 1 or 2 (depending on rack size) 3m InfiniBand cables
- 1 blue Ethernet cable

- 1 red Ethernet cable
- 1 black Ethernet cable

Two sets of two keys are provided for the rack door and side panel locks.

About Tape Backup Infrastructure

Tape backup is an optional feature. You can choose between Oracle and third-party tape systems.

Oracle's Recovery Appliance supports the installation and configuration of Oracle tape systems. You can connect Recovery Appliance to the storage area network (SAN) in your data center for backups to an Oracle tape system.

Oracle Recommended Stack

Recovery Appliance is engineered to work with industry leading tape infrastructure components, including those described in Table 1-18.

Table 1-18 Oracle Tape System Components

Component	Models	Notes
Oracle Tape Libraries	SL150, SL3000, SL8500	ACSLS, ACSLS Logical Libraries, FC-SCSI
Oracle Tape Drives	T10000D, T10000C, LTO	All interfaces supported by the recommended Oracle libraries
Automated Cartridge System Library Software (ACSLS)	Automated Cartridge System Library Software (ACSLS) 8.x	API, Logical Libraries
Tape Encryption	Oracle Key Manager (OKM) 2.3	Recommended Key Management for enterprise key management. Operates seamlessly with Recovery Appliance

Recovery Appliance supports Oracle Secure Backup tape management software out of the box. It provides the following benefits:

- Oracle Secure Backup is deployed and configured as part of the Recovery Appliance software installation.
- Upgrades are included with Recovery Appliance software upgrades.
- Tape backups automatically protect the Recovery Appliance metadata.
- Oracle Secure Backup provides disaster recovery automation from tape.
- Oracle Secure Backup copies data from its storage location directly to tape.
- Oracle Secure Backup provides superior local tape performance by using only the InfiniBand and fiber channel networks. Tape backups do not impact your IP network.

Third-Party Tape Systems

Out of the box, Recovery Appliance supports the use of any third-party (non Oracle) media management software that is compatible with the Recovery Appliance.

You can install the *client* software component of a third-party media management product, such as Symantec NetBackup or HP Data Protector, on the compute servers. You can then configure Recovery Appliance as a client over the 10 GbE network.

Under the terms of the Recovery Appliance license, you may not install the media manager *server* component because only Oracle Secure Backup can back up Recovery Appliance to the tape devices directly (or via SAN) attached to the fibre channel interfaces of the compute nodes. Any non Oracle media manager software must be configured only as a network client.

Note the following:

- Recovery Appliance software upgrades might not preserve the client software or its configuration.
- Third-party media managers do not provide Recovery Appliance metadata protection support.
- If you configure third-party media managers with Recovery Appliance, carefully validate your configuration to ensure that backups and restores work as expected and leverage all nodes in the Recovery Appliance RAC cluster.

Restrictions on Modifying Recovery Appliance

The following restrictions apply to software and hardware modifications to Recovery Appliance. Violating these restrictions can result in loss of warranty and support.

- Customers may not modify any internal software component or configuration on a Recovery Appliance compute server except as indicated in a My Oracle Support note or as explicitly directed by an Oracle support analyst. This includes actions such as modifying or adding listeners to the internal database, modifying the operating system or network configuration files, installing custom scripts, creating NFS mounts for any purpose except for polling a backup location, and creating additional database instances.
- Although Oracle discourages installation of third-party software products on the appliance, customers may install them under the conditions stated in My Oracle Support note 2014361.1.
- Customers may install the client software component of a third-party media management product in "LAN backup agent" mode on the compute servers. Customers may not install the media manager server software component, because only Oracle Secure Backup supports Recovery Appliance backups to tape over fibre channel connections.
- Customers may not load additional software or alter the installed software on a storage server.
- Customers may not install Oracle Enterprise Manager Ops Center agents. However, Oracle Enterprise Manager Cloud Control Management agents can be installed on the compute servers to monitor the system, as part of the Recovery Appliance software installation process.
- Recovery Appliance hardware cannot be modified or customized. There are two exceptions. The allowed hardware modifications to Recovery Appliance are as follows:

  - Modification to the administrative 48-port Cisco Gigabit Ethernet switch included with Recovery Appliance. Customers may choose to do the following:
    * Replace the Gigabit Ethernet switch, at their own expense, with an equivalent 1U 48-port Gigabit Ethernet switch that conforms to their internal data center network standards. This replacement must be performed by the customer, at their expense and labor, after delivery of Recovery Appliance. If the customer chooses to make this change, then Oracle cannot make or assist with this change given the numerous possible scenarios involved, and it is not included as part of the standard installation. The customer must supply the replacement hardware, and make or arrange for this change through other means.
    * Remove the CAT5 cables connected to the Cisco Ethernet switch, and connect them to the customer's network through an external switch or patch panel. The customer must perform these changes at their expense and labor. In this case, the Cisco Ethernet switch in the rack can be turned off and disconnected from the data center network.
  - Addition of client access network switches in Recovery Appliance. The following restrictions apply for the network switches:
    * A maximum of two client access network switches may be installed in the rack.
    * The switch must be 1 rack unit (RU) in height.
    * Air flow must go from the front of the rack to the back of the rack.
    * The switch must be installed in rack slot U41 or U42.
    * Power consumption must be less than 400 watts for each switch.
    * Power over Ethernet (PoE) functionality, if applicable, should not be used on these switches.
- Customers can update the firmware of the other components.
  - Customers can update the IOS and firmware versions on the Cisco Gigabit Ethernet switch to meet their data center requirements.
  - Customers can update the firmware of the components of the compute servers only by applying a patch bundle provided by Oracle.
  - Customers can update the firmware of the InfiniBand switches provided they comply with the validated versions documented in My Oracle Support note 1927416.1, and its related notes.
- Customers cannot connect USB devices that draw more than 100 mA of power to the compute servers.

2

Preparing Your Networks for Recovery Appliance

This chapter describes the network requirements for Recovery Appliance, so that you can prepare your data center for installation.

This chapter contains the following sections:

Overview of Network Requirements

In addition to the compute and storage servers, Recovery Appliance includes equipment to connect the system to your network. The network connections allow clients to connect to the compute servers and also enables remote system administration.

Use the information in this section in conjunction with Oracle Exadata Deployment Assistant (OEDA) to configure your Recovery Appliance environment.

To deploy Recovery Appliance ensure that you meet the minimum network requirements. Recovery Appliance requires a minimum of three networks, and there are interfaces available for additional networks. Each network must be on a separate and distinct subnet. The network descriptions are as follows:

- **Administration Network**: Also known as the management network, this required network connects to your existing management network infrastructure, and is used for administrative work on all components of Recovery Appliance. By default, the administration network connects the compute servers, storage servers, server Integrated Lights Out Manager (ILOM) interfaces, and RDMA Network Fabric switches to the Cisco Management Switch in the rack. One uplink is required from the Cisco Management Switch to your management network.

  Each compute server and storage server has two network interfaces for administration. One interface provides management access to the operating system through a dedicated Ethernet port. The other network interface is dedicated to ILOM. By default, Recovery Appliance is delivered with both interfaces connected to the Cisco Management Switch. Cabling or configuration changes to these interfaces is not permitted, except that the ILOM interfaces can be connected to a dedicated ILOM network, which is separate from the administration network. The administration network interfaces on the compute servers should not be used for client or application network traffic.

> **Notes:**
>
> – Separate uplinks to your management network are also recommended for remote monitoring of each power distribution unit (PDU). This configuration enables you to easily differentiate between system outages caused by PDU failure as opposed to failure of the Cisco Management Switch.
>
> – A properly secured configuration requires full isolation of the administration network from all other networks.

- **Ingest Network**: This required network connects the protected Oracle Database servers to Recovery Appliance for backup within the same data center. Also known as a backup network, this high-speed, private Ethernet network must be designed to support the transfer of large volumes of data. Recovery Appliance connects to this network using two 10/25 GB connections to each of the two compute servers in the rack. You can configure the two connections as active/passive (redundant) or active/active.

  The compute servers support channel bonding to provide higher bandwidth and availability.

  Single client access name (SCAN) supports failover between the two compute servers in the Recovery Appliance. In an installation with multiple Recovery Appliance racks configured as a cluster, virtual IP (VIP) addresses support failover among the racks. The protected database systems can resolve the host names to dynamically assigned addresses.

  Third-party tape hardware and software also uses the ingest network.

- **Private Network**: Also known as the RDMA Network Fabric, or the interconnect, this network connects the compute servers and storage servers. Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on the Oracle Exadata Storage Servers. The private network is automatically configured during installation. It is non-routable, fully contained in Recovery Appliance, and does not connect to your existing networks.

  Starting with Recovery Appliance X8M, the private network uses RDMA over Converged Ethernet (RoCE).

  Previously, the private network was built using InfiniBand technology. RoCE Network Fabric uses different switches and cables from those used by InfiniBand Network Fabric.

- **Replication network**: The optional replication network uses available ports not used by the administration and the ingest network. It connects the local Recovery Appliance (the **upstream** appliance) with a remote Recovery Appliance (the **downstream** appliance). Oracle recommends a broadband, encrypted network, instead of an insecure public network, wherever possible.

  Recovery Appliance supports the following configurations between the upstream and downstream appliances:

> **Note:**
>
> A downstream Recovery Appliance or a tape library can reside in the local data center. The replication network is not used in a local configuration.
>
> The replication network must not be used for the purpose of ingesting backups.

- **Fiber Channel SAN network**: If you are using Oracle Secure Backup, then you can back up Recovery Appliance to the storage area network (SAN) in your data center for backups to tape. The network connections depend on whether you have an Oracle tape solution or use third-party hardware.

Ingest and replication networks can be configured active/passive or active/active bonding.

- Active / Passive Bonding - BONDING_OPTS="mode=active-backup miimon=100 downdelay=2000 updelay=5000 num_grat_arp=100"
- Active / Active Bonding - BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200 lacp_rate=1 xmit_hash_policy=layer3+4"

Ingest can be Active/Active with Replication Active/Passive, or vice-versa. Or both can have the same bonding.

> **See Also:**
>
> - *Oracle Clusterware Administration and Deployment Guide* for a discussion of SCANs and VIPs in network configurations for Oracle Database.
> - "Connecting Recovery Appliance to a Tape Library" for information about how a fibre channel SAN network is configured for backups to tape in a Recovery Appliance environment.

About the X9M and X8M Network Components and Interfaces

Each compute server in the X9M and X8M configuration consists of the following network components and interfaces:

- Ethernet ingest and replication network connectivity, :
  - 2 x Dual 10/25G Network Cards, or
  - 2 x Quad 10G Network Cards, or
  - 1 x Quad 10G Network Card and 1 x Dual 10/25G Network Card
- 2 x QSFP2B RoCE Fabric Ports
- 1 Ethernet port for `Serial MGT` remote management
- 1 Ethernet port for `ILOM MGT` (Oracle Integrated Lights Out Manager) remote management
- 1 Ethernet port for `HOST MGT` remote management
- Optional: Dual 32G HBA (Tape) card that can be field installed. This slot is not available for other network cards.

Figure 2-1 X9M Backplane External Network Connectivity

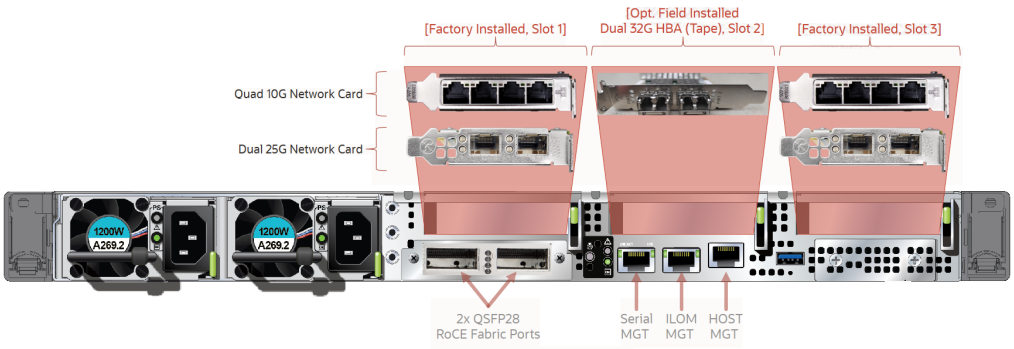

- Maximum of two 10G (or 25G) ports for ingest and two 10G (or 25G) ports for replication, per compute server
- Maximum of four 10G (or 25G) for each network in bonded LACP configuration, per rack
- Replication can be used as separate ingest network (MOS Note 2126047.1)
- VLAN tagging supported on ingest network (MOS Note 2047411.1)
- Slot 2 cannot be used for additional network card.

Figure 2-2 X8M Backplane External Network Connectivity

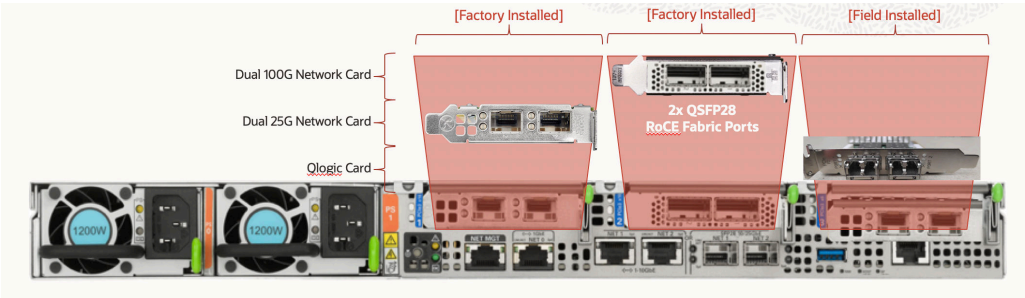

- Maximum of two 25G ports for ingest and two 25G ports for replication, per compute server

About the X8 and X7 Network Components and Interfaces

Each compute server in the X8-2 and X7 configuration consists of the following network components and interfaces:

- Ethernet ingest and replication network connectivity
  - On-board: 2 x 10 Gb copper Ethernet (eth1)
  - On-board: 2 x 10/25 Gb optical Ethernet Ports (eth2)
  - PCIe card: 2 x 10/25 Gb optical Ethernet Ports (eth3 and eth4)
- 1 dual-port 4X QDR (40 Gbps) InfiniBand Host Channel Adapter (HCA) (IB0 and IB1)

- 1 Ethernet port for Oracle Integrated Lights Out Manager (ILOM) remote management
- 1 dual-port 32 GB FC Converged Network Adapter (CNA) FC ports 0 and 1.

> **Note:**
>
> The corresponding SFP modules that work with the 10/25 GbE PCIe 2.0 network cards are purchased separately.

Figure 2-3 X8-2 and X7 Backplane External Network Connectivity

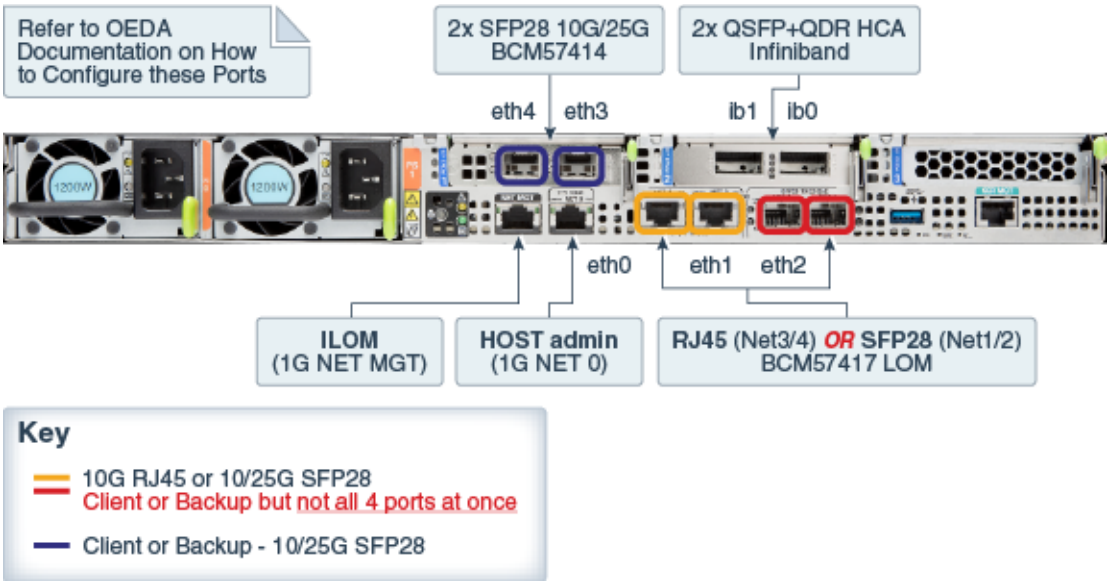

Given that the base rack has two (2) compute servers, the maximum for ingest is 2 x 10 Gb or 2 x 25 Gb Ethernet ports, while the maximum for replication is 2 x 10 Gb or 2 x 25 Gb Ethernet ports. The following are valid combinations of the options.

- 2 x 10Gb on-board copper (ingest) + 2 x 10/25Gb PCIe card optical (replication)
- 2 x 10/25Gb PCIe card optical (ingest) + 2 x 10Gb on-board copper (replication)
- 2 x 10/25Gb PCIe card optical (ingest) + 2 x 10/25Gb on-board optical (replication)
- 2 x 10/25Gb on-board optical (ingest) + 2 x 10/25Gb PCIe card optical (replication)

> **Note:**
>
> If ingest and replication traffic is desired to be configured on the **same network,** define the required network interface in OEDA in the ingest network section, and leave the replication network section blank. With this setup, Recovery Appliance will use the ingest network for replication traffic.

Each storage server consists of the following network components and interfaces:

- 1 embedded Gigabit Ethernet port (NET0)
- 1 dual-port 4X QDR (40 Gbps) InfiniBand Host Channel Adapter (HCA) (IB0 and IB1)
- 1 Ethernet port for Oracle Integrated Lights Out Manager remote management (Oracle ILOM)

Additional configuration, such as defining multiple virtual local area networks (VLANs) for the management (NET0 and/or ILOM) interfaces or enabling routing, might be required for the switch to operate properly in your environment and is beyond the scope of the installation service. If additional configuration is needed, then your network administrator must perform the necessary configuration steps during installation of Recovery Appliance.

Example of Network Connections for Recovery Appliance

Figure 2-4 shows the network cabling of a sample configuration. Two Recovery Appliance racks are installed in separate data centers. The protected Oracle databases are connected to the upstream Recovery Appliance over the ingest network. The upstream Recovery Appliance is connected to the downstream Recovery Appliance over the replication network. Both racks are configured to use an Oracle tape solution.

Figure 2-4 Network Diagram for Recovery Appliance

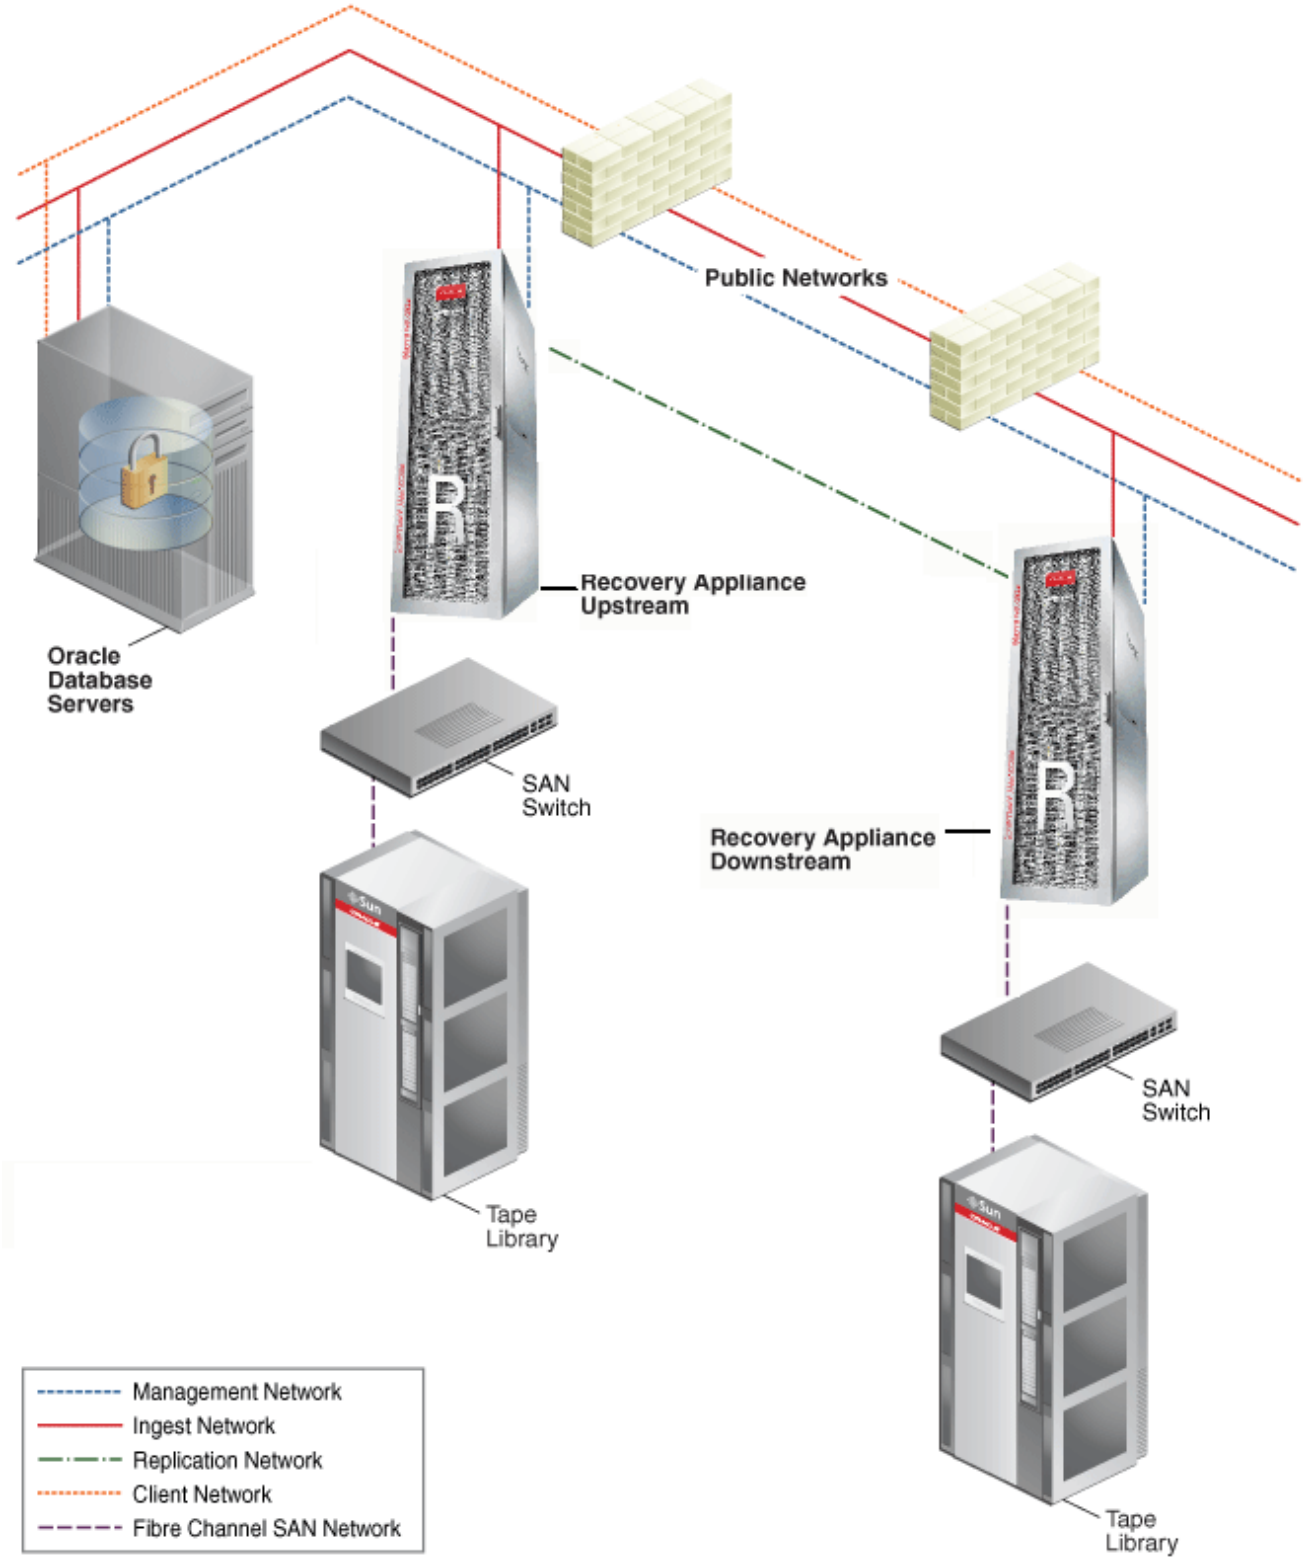

Connecting Recovery Appliance Rack Components to the Networks

#unique_75/unique_75_Connect_42_CHDJGBEC shows the network connections to components of Recovery Appliance rack.

The management network connects through the Ethernet switch to the compute servers, the storage servers, and the RDMA Network Fabric switches. The management network connects directly to the PDUs.

The ingest network, the optional replication network, and the optional fiber channel SAN network connect to the two compute servers.

The RDMA Network Fabric network connects the switches to the compute servers and the storage servers.

Figure 2-5 Network Connections to the Recovery Appliance Rack Components

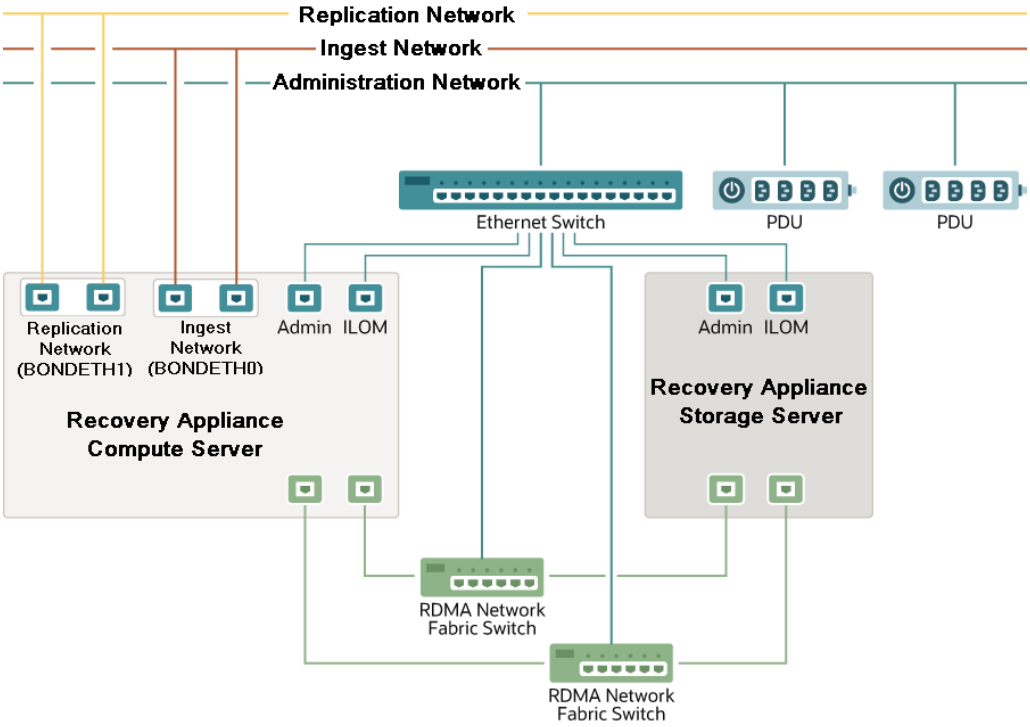

Connecting Recovery Appliance to a Tape Library

The network connections between Recovery Appliance and an optional tape library depend on whether you are using Oracle or third-party tape management system. See "About Tape Backup Infrastructure" for the differences in support provided by Recovery Appliance.

Oracle Recommended Stack

When you use the Oracle compatible tape solution, a fiber channel adapter is installed in each compute server to provide a connection to the fiber channel storage area network (SAN). Tape backups are isolated on this network, and thus do not interfere

with the performance of the other networks. Figure 2-6 provides an overview of the network connections when using an Oracle tape system.

Figure 2-6 Recovery Appliance Connection to an Oracle Tape System

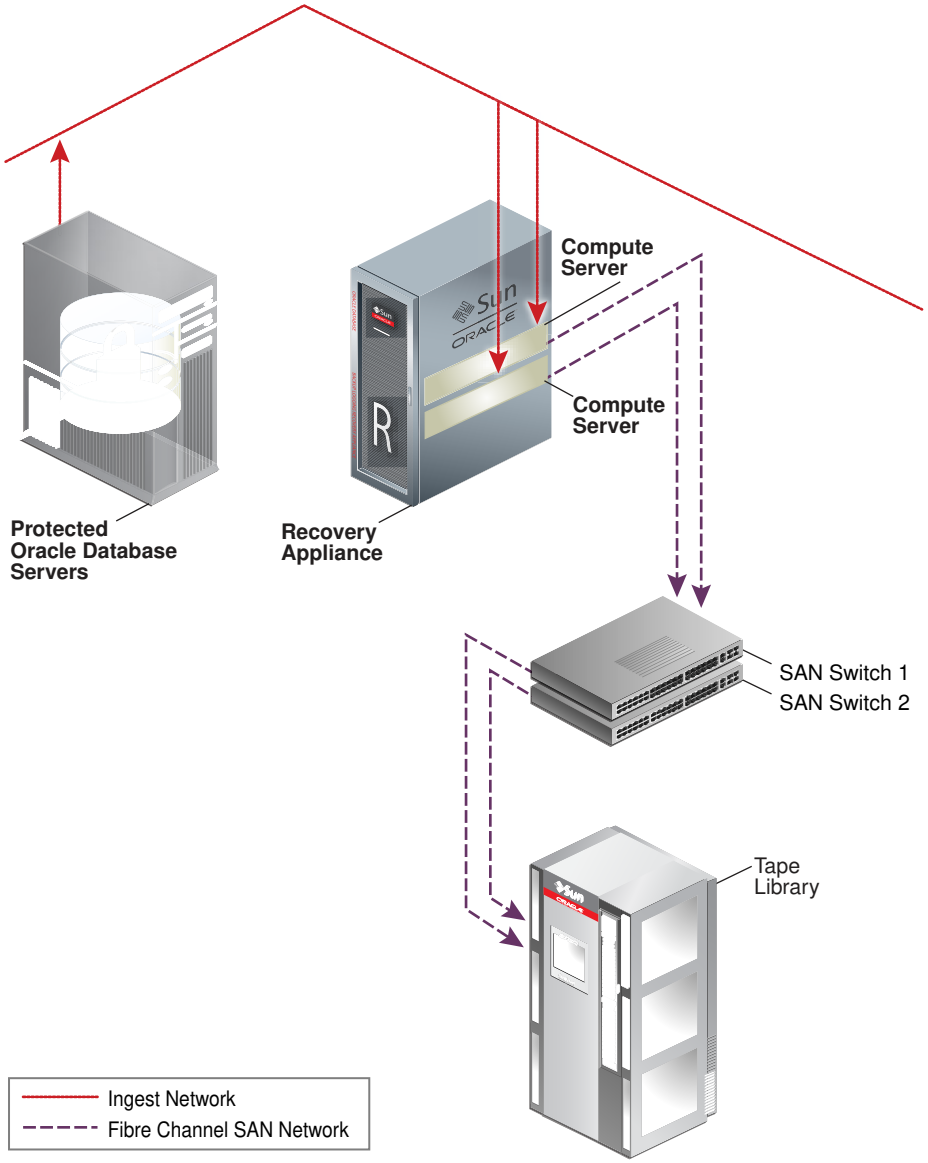

Third-Party Tape Systems

When you use a third-party tape system, the backups to tape use the 10 Gb ingest network. This is the same network that the local protected databases use to backup to Recovery Appliance. Figure 2-7 provides an overview of the network connections when using a third-party tape system.

Figure 2-7 Recovery Appliance Connection to a Third-Party Tape System

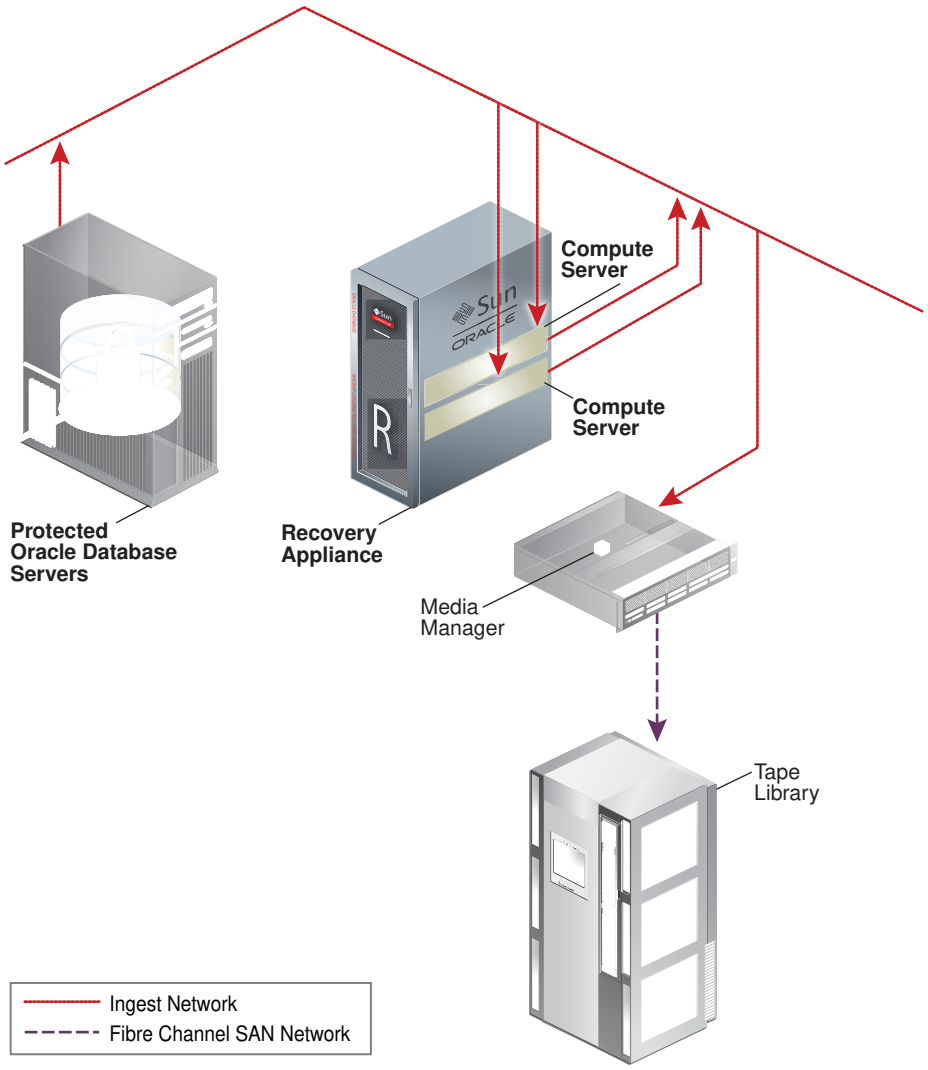

Using Network VLAN Tagging with Recovery Appliance

The Recovery Appliance supports VLAN port tagging only on the ingest network. You configure VLAN port tagging after you complete the Recovery Appliance installation.

If applicable, ensure that you also set the Access VLAN on the network switches, including on the Cisco switch that is included in the Recovery Appliance rack for the management network.

> **See Also:**
>
> "Installing the Software on Recovery Appliance" for instructions on when and how to configure VLAN tagging

Registering Recovery Appliance in the Domain Name System

Before receiving your Recovery Appliance rack, use Oracle Exadata Deployment Assistant. The assistant generates a file to be used when setting up the system. The host names and IP addresses specified in the assistant-generated file must be registered in Domain Name System (DNS) before the initial configuration. In addition, all public addresses, single client access name (SCAN) addresses, and VIP addresses must be registered in DNS before installation.

The assistant-generated file defines the SCAN as a single name with three IP addresses on the client access network. The three SCAN addresses provide service access for clients to Recovery Appliance. Configure DNS for round robin resolution for the SCAN name to these three SCAN addresses.

All addresses registered in DNS must be configured for both forward resolution and reverse resolution. Reverse resolution must be forward confirmed (forward-confirmed reverse DNS) such that both the forward and reverse DNS entries match each other.

See Also:

- Using the Deployment Assistant
- *Oracle Grid Infrastructure Installation Guide for Linux* for additional information about SCAN addresses
- Your DNS vendor documentation for additional information about configuring round-robin name resolution

Factory IP Address Settings

Recovery Appliance has default IP addresses set at the factory:

- Gateway: 192.168.1.254 in all devices as required
- Subnet Mask: 255.255.252.0 in all devices as required
- IP Address Range: 192.168.1.1 to 192.168.1.203

Before connecting Recovery Appliance to the network, ensure that these IP addresses do not conflict with other addresses on the network. The `checkip.sh` script checks for conflicts. Oracle recommends running the script before connecting the network to avoid problems, even when a check was performed before Recovery Appliance was delivered. See "Installing the Software on Recovery Appliance" for additional information about the `checkip.sh` script.

Table 2-1 lists the factory IP addresses for a Recovery Appliance full rack.

Table 2-1 Factory IP Addresses for Recovery Appliance

Rack Unit	Component	Management Network Addresses	InfiniBand Active Bonded IP Addresses	Oracle ILOM IP Addresses
U41	Storage server	192.168.1.23	192.168.10.45	192.168.1.123

Table 2-1 (Cont.) Factory IP Addresses for Recovery Appliance

Rack Unit	Component	Management Network Addresses	InfiniBand Active Bonded IP Addresses	Oracle ILOM IP Addresses
U39	Storage server	192.168.1.22	192.168.1.43	192.168.1.122
U37	Storage server	192.168.1.21	192.168.10.41	192.168.1.121
U35	Storage server	192.168.1.20	192.168.10.39	192.168.1.120
U33	Storage server	192.168.1.19	192.168.10.37	192.168.1.119
U31	Storage server	192.168.1.18	192.168.10.35	192.168.1.118
U29	Storage server	192.168.1.17	192.168.10.33	192.168.1.117
U27	Storage server	192.168.1.16	192.168.10.31	192.168.1.116
U25	Storage server	192.168.1.14	192.168.10.27	192.168.1.114
U23	Storage server	192.168.1.12	192.168.10.23	192.168.1.112
U22	RDMA Network Fabric switch	Not applicable	Not applicable	192.168.1.203
U21	Ethernet switch	Not applicable	Not applicable	192.168.1.200
U20	RDMA Network Fabric switch	Not applicable	Not applicable	192.168.1.202
U18	Storage server	192.168.1.10	192.168.10.19	192.168.1.110
U17	Compute server	192.168.1.9	192.168.10.17	192.168.1.109
U16	Compute server	192.168.1.8	192.168.10.15	192.168.1.108
U14	Storage server	192.168.1.7	192.168.10.13	192.168.1.107
U12	Storage server	192.168.1.6	192.168.10.11	192.168.1.106
U10	Storage server	192.168.1.5	192.168.10.9	192.168.1.105
U08	Storage server	192.168.1.4	192.168.10.7	192.168.1.104
U06	Storage server	192.168.1.3	192.168.10.5	192.168.1.103
U04	Storage server	192.168.1.2	192.168.10.3	192.168.1.102
U02	Storage server	192.168.1.1	192.168.10.1	192.168.1.101

Port Assignments When Using a Firewall

When network communication between Recovery Appliance and other components requires access through a firewall, you must open ports used by the Recovery Appliance services.

> **Note:**
>
> A firewall may not be used between components of the Recovery Appliance.

Table 2-2 lists the ports used by services on Recovery Appliance. Review the list and open the necessary ports. All ports are on the management network, unless otherwise noted.

Table 2-2 Open Ports for the Firewall

Source	Target	Protocol	Port	Application
NA	Database management	SSH over TCP	22	SSH
NA	Compute servers, storage servers, and InfiniBand ILOMs	SSH over TCP	22	SSH
NA	Storage management	SSH over TCP	22	SSH
Storage servers	email server	SMTP	25 465 if using SSL	SMTP (Simple Mail Transfer Protocol)
Compute servers, storage servers, and InfiniBand ILOMs	NA	TFTP over UDP	69	Outgoing TFTP (Trivial File Transfer Protocol)
NA	Compute servers, storage servers, and InfiniBand ILOMs	HTTP over TCP	80	Web (user configurable)
NA	PDU	HTTP over TCP	80	Browser interface
Database management	NA	NTP over UDP	123	Outgoing Network Time Protocol (NTP)
Compute servers, storage servers, and InfiniBand ILOMs	NA	NTP over UDP	123	Outgoing NTP
Storage management	NA	NTP over UDP	123	Outgoing NTP
ASR Manager	ASR asset	SNMP (get)	161	FMA enrichment for additional diagnostic information
NA	Compute servers, storage servers, and InfiniBand ILOMs	SNMP over UDP	161	SNMP (Simple Network Management Protocol) (user configurable)
NA	PDU	SNMP over UDP	161	SNMP (user configurable)
Storage servers	SNMP subscriber such as Oracle Enterprise Manager Cloud Control or an SNMP manager	SNMP	162	SNMP version 1 (SNMPv1) outgoing traps (user-configurable)
Compute servers and storage server ILOMs	ASR Manager	SNMP	162	Telemetry messages sent to ASR Manager
Compute servers, storage servers, and InfiniBand ILOMs	NA	IPMI over UDP	162	Outgoing IPMI (Intelligent Platform Management Interface) Platform Event Trap (PET)

Table 2-2 (Cont.) Open Ports for the Firewall

Source	Target	Protocol	Port	Application
PDU	NA	SNMP over UDP	162	Outgoing SNMPv2 traps
NA	Compute servers, storage servers, and InfiniBand ILOMs	LDAP over UDP/TCP	389	Outgoing LDAP (Lightweight Directory Access Protocol) (user configurable)
ASR Manager	ASR back end	HTTPS	443	Telemetry messages sent to ASR back end
NA	Compute servers, storage servers, and InfiniBand ILOMs	HTTPS over TCP	443	Web (user configurable)
NA	PDU	HTTPS over TCP	443	Browser interface
Compute servers, storage servers, and InfiniBand ILOMs	NA	Syslog over UDP	514	Outgoing Syslog
PDU	NA	Syslog over UDP	514	Outgoing Syslog
Compute servers, storage servers, and InfiniBand ILOMs	NA	DHCP over UDP	546	client DHCP (Dynamic Host Configuration Protocol)
PDU	NA	DHCP over UDP	546	DHCP (Dynamic Host Configuration Protocol) client
NA	Compute servers, storage servers, and InfiniBand ILOMs	IPMI over UDP	623	IPMI (Intelligent Platform Management Interface)
Oracle Enterprise Manager Cloud Control	NA	TCP	1159	Oracle Enterprise Manager Cloud Control HTTPS upload port
Oracle Enterprise Manager Cloud Control	NA	TCP	1159	Oracle Enterprise Manager Cloud Control HTTPS upload port
NA	Database data	SQL*Net over TCP	1521	Database listener
Protected database	Recovery Appliance	SQL*Net over TCP	1521 (ingest network)	RMAN backup and restore
Upstream Recovery Appliance	Downstream Recovery Appliance	SQL*Net over TCP	1522 (replication network)	Recovery Appliance Replication
Compute servers, storage servers, and InfiniBand ILOMs	NA	RADIUS over UDP	1812	Outgoing RADIUS (Remote Authentication Dial In User Service) (user configurable)

Table 2-2 (Cont.) Open Ports for the Firewall

Source	Target	Protocol	Port	Application
Oracle Enterprise Manager Grid Control	NA	TCP	4889	Oracle Enterprise Manager Cloud Control HTTP upload port
Oracle Enterprise Manager Grid Control	NA	TCP	4889	Oracle Enterprise Manager Cloud Control HTTP upload port
NA	Compute server and storage server ILOMs	TCP	5120	ILOM remote console: CD
NA	Compute server and storage server ILOMs	TCP	5121	ILOM remote console: keyboard and mouse
NA	Compute server and storage server ILOMs	TCP	5123	ILOM remote console: diskette
NA	Compute server and storage server ILOMs	TCP	5555	ILOM remote console: encryption
NA	Compute server and storage server ILOMs	TCP	5556	ILOM remote console: authentication
ASR Manager	Compute server and storage server ILOMs	HTTP	6481	Service tags listener for asset activation
NA	Compute server and storage server ILOMs	TCP	6481	ILOM remote console: `servicetag` daemon
NA	Compute server and storage server ILOMs	TCP	7578	ILOM remote console: video
NA	Compute server and storage server ILOMs	TCP	7579	ILOM remote console: serial
NA	Compute servers	TCP	7777	Oracle Enterprise Manager Grid Control HTTP console port
NA	Storage servers	TCP	7777	Oracle Enterprise Manager Grid Control HTTP console port
NA	Compute servers	TCP	7799	Oracle Enterprise Manager Grid Control HTTPS console port
NA	Storage servers	TCP	7799	Oracle Enterprise Manager Grid Control HTTPS console port

Table 2-2 (Cont.) Open Ports for the Firewall

Source	Target	Protocol	Port	Application
Protected database	Recovery Appliance	HTTP	8001 (ingest network)	RMAN backup and restore
Upstream Recovery Appliance	Downstream Recovery Appliance	HTTP	8001 (replication network)	Recovery Appliance Replication

3

Setting Up Auto Service Request

This chapter explains how to install and configure Auto Service Request for Recovery Appliance. It contains these sections:

Understanding Auto Service Request

Auto Service Request (ASR) is designed to automatically open service requests when specific Recovery Appliance hardware faults occur. ASR detects faults in the most common server components, such as disks, fans, and power supplies, and automatically opens a service request when a fault occurs. ASR monitors only server components and does not detect all possible faults.

ASR is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts, within the customer data center. It is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.

When ASR detects a hardware problem, ASR Manager submits a service request to Oracle Support Services. In many cases, Oracle Support Services can begin work on resolving the issue before the administrator is even aware the problem exists.

An email message is sent to both the My Oracle Support email account and the technical contact for Recovery Appliance to notify them of the creation of the service request.

A service request might not be filed automatically on some occasions. This can happen because of the unreliable nature of the SNMP protocol or a loss of connectivity to ASR Manager. Oracle recommends that you continue to monitor your systems for faults and call Oracle Support Services if you do not receive notice that a service request was filed automatically.

See Also:

- Oracle Auto Service Request web page at

  http://www.oracle.com/technetwork/systems/asr/overview/index.html

- Oracle Auto Service Request user documentation at

  http://www.oracle.com/technetwork/systems/asr/documentation/index.html

Getting Ready to Install ASR

Before installing ASR, complete these prerequisites:

1. Create a My Oracle Support account at http://support.oracle.com.
2. Ensure that the following are set up correctly in My Oracle Support:
   - Oracle Premier Support for Systems, Oracle Premier Support for Operating Systems, or Oracle Limited Warranty
   - The technical contact person at the customer site who is responsible for Recovery Appliance
   - A valid shipping address at the customer site for the Recovery Appliance parts
3. Identify and designate a system to host ASR Manager.

   ASR Manager must be installed on a server that has connectivity to Recovery Appliance and an outbound Internet connection using HTTPS or an HTTPS proxy. To submit a service request (SR), the server must be able to access the Internet.
4. ASR Manager uses the following ports for normal operations. Ensure that Recovery Appliance and ASR Manager server ports are open:
   - All Recovery Appliance servers must have port 6481 open for HTTP-based asset activation requests that they receive from ASR Manager.
   - The server running ASR Manager must have port 162 open for SNMP-based telemetry messages that the Recovery Appliance servers send to it.
   - The Oracle `transport.oracle.com` web server has port 443 open for HTTPS-based telemetry messages that ASR Manager sends to it.

   See Figure 3-1 for the port locations.
5. Ensure that the designated system conforms to the "Hardware and Network Configuration Recommendations for ASR." See the Oracle ASR website:

   http://docs.oracle.com/cd/E37710_01/nav/products.htm
6. Confirm that Java Development Kit 6 (JDK 1.6.0_04 or later) is running on the designated ASR Manager system:

   ```
   java -version
   ```

   If necessary, download and install the latest version of JDK from the Java SE Downloads website:

http://www.oracle.com/technetwork/java/javase/downloads/index.html

7. Obtain `root` access to the designated ASR Manager host.
8. Identify and verify connectivity to Recovery Appliance.
9. Verify connectivity to the Internet using HTTPS.

Figure 3-1 shows the network connections between ASR and Recovery Appliance.

Figure 3-1 Auto Service Request Network Connections

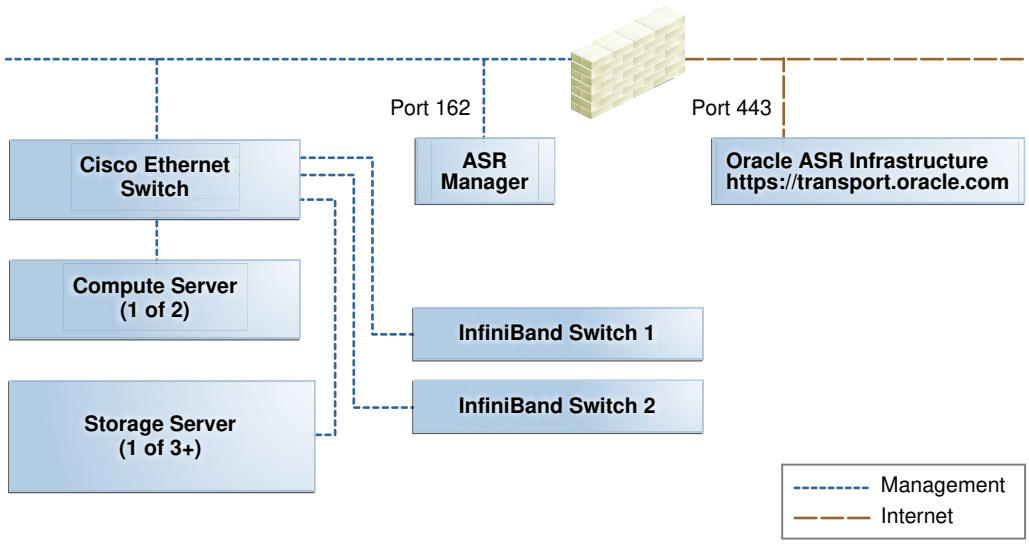

Installing ASR Manager

To install ASR Manager, download the current version from My Oracle Support Doc ID 1185493.1. Then follow the instructions in the *Oracle Auto Service Request Installation and Operations Guide* at

http://www.oracle.com/technetwork/systems/asr/documentation/index.html

Verifying the ASR Manager Installation

Perform these checks as `root` on ASR Manager to ensure that it is installed properly.

- Ensure that the `asr` command is in the search path of the `root` user.
- Verify that ASR Manager 3.5 or later is running:

  ```
  asr show_rules_version
  ```

- Check the registration status:

  ```
  asr show_reg_status
  ```

- Test the connection by sending a test message to the transport server.

  ```
  asr test_connection
  ```

- Verify the ASR assets:

1. In ASR Manager, verify that ASR is activated:

   ```
   asr list_asset -i asset_ip
   ```

   In the preceding command, *asset_ip* is the IP address of a server or an Oracle ILOM. To list all assets, enter this command:

   ```
   asr list_asset
   ```

2. If no assets are listed, then verify that ASR is configured on Recovery Appliance.

> **Note:**
>
> If an IP address or host name changes, then you must deactivate and reactivate the asset.

About the Trap Destinations on Recovery Appliance

ASR is an optional software component of Recovery Appliance. When configuring the software in Oracle Exadata Deployment Assistant, you must complete the page about ASR. Then during the software installation, Recovery Appliance configures the servers to trap the component fault indicators. See "#unique_86".

The following are examples of Simple Network Management Protocol (SNMP) traps sent to ASR Manager for disk failures.

shows the SNMP trap for an storage server disk failure. The hardware alert code (HALRT-02001) is underlined.

shows the SNMP trap from a compute server disk failure. The hardware alert code (HALRT-02007)is underlined.

Example 3-1 Example of a Storage Server SNMP Trap

```
2011-09-07 10:59:54 server1.example.com [UDP: [192.85.884.156]:61945]:
RFC1213-MIB::sysUpTime.0 = Timeticks: (52455631) 6 days, 1:42:36.31
SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: SUN-HW-TRAP-MIB::sunHwTrapHardDriveFault
SUN-HW-TRAP-MIB::sunHwTrapSystemIdentifier = STRING: Sun Oracle Database Machine
1007AK215C
SUN-HW-TRAP-MIB::sunHwTrapChassisId = STRING: 0921XFG004
SUN-HW-TRAP-MIB::sunHwTrapProductName = STRING: SUN FIRE X4270 M2 SERVER
SUN-HW-TRAP-MIB::sunHwTrapSuspectComponentName = STRING: SEAGATE
ST32000SSSUN2.0T;
Slot: 0SUN-HW-TRAP-MIB::sunHwTrapFaultClass = STRING: NULL
SUN-HW-TRAP-MIB::sunHwTrapFaultCertainty = INTEGER: 0
SUN-HW-TRAP-MIB::sunHwTrapFaultMessageID = STRING: HALRT-02001
SUN-HW-TRAP-MIB::sunHwTrapFaultUUID = STRING:
acb0a175-70b8-435f-9622-38a9a55ee8d3
SUN-HW-TRAP-MIB::sunHwTrapAssocObjectId = OID: SNMPv2-SMI::zeroDotZero
SUN-HW-TRAP-MIB::sunHwTrapAdditionalInfo = STRING: Exadata Storage Server:
cellname  Disk Serial Number:   E06S8K
server1.example.com failure trap.
```

Example 3-2 Example of Oracle Database Server SNMP Trap

```
2011-09-09 10:59:54 dbserv01.example.com [UDP: [192.22.645.342]:61945]:
RFC1213-MIB::sysUpTime.0 = Timeticks: (52455631) 6 days, 1:42:36.31
```

```
SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: SUN-HW-TRAP-MIB::sunHwTrapHardDriveFault
SUN-HW-TRAP-MIB::sunHwTrapSystemIdentifier = STRING: Sun Oracle Database Machine
1007AK215C
SUN-HW-TRAP-MIB::sunHwTrapChassisId = STRING: 0921XFG004
SUN-HW-TRAP-MIB::sunHwTrapProductName = STRING: SUN FIRE X4170 M2 SERVER
SUN-HW-TRAP-MIB::sunHwTrapSuspectComponentName = STRING: HITACHI H103030SCSUN300G
Slot: 0SUN-HW-TRAP-MIB::sunHwTrapFaultClass = STRING: NULL
SUN-HW-TRAP-MIB::sunHwTrapFaultCertainty = INTEGER: 0
SUN-HW-TRAP-MIB::sunHwTrapFaultMessageID = STRING: HALRT-02007
SUN-HW-TRAP-MIB::sunHwTrapFaultUUID = STRING: acb0a175-70b8-435f-9622-38a9a55ee8d3
SUN-HW-TRAP-MIB::sunHwTrapAssocObjectId = OID: SNMPv2-SMI::zeroDotZero
SUN-HW-TRAP-MIB::sunHwTrapAdditionalInfo = STRING: Exadata Database Server: db03
Disk Serial Number: HITACHI H103030SCSUN300GA2A81019GGDE5E
dbserv01.example.com failure trap.
```

Troubleshooting ASR

For troubleshooting procedures for the ASR software, see Chapter 5 of the *Oracle ASR Installation and Operations Guide* at:

http://www.oracle.com/technetwork/systems/asr/documentation/index.html

If you continue to have issues, contact ASR support.

4
Using the Deployment Assistant

This chapter describes how to use the Deployment Assistant. The information that you enter in it is used to create the Recovery Appliance configuration files. These files work with installation utility to automate the installation process.

Overview of Oracle Exadata Deployment Assistant

OEDA is used to create configuration files and then implement the configuration on your engineered system.

OEDA is used for Oracle Exadata Database Machine, Oracle Zero Data Loss Recovery Appliance, and Oracle SuperCluster. Starting with Oracle Exadata System Software release 19.1.0, there is a new Web-based interface for OEDA. However, the overall process, described here, is the same regardless of whether you use the Web-based or Java-based user interface.

- Before your engineered system arrives, do the following:
  1. Work with your network and database administrators to evaluate the current network settings, such as current IP address use and network configuration.
  2. Define the settings for the rack, such as network configuration and backup method.
  3. Download the latest version of OEDA from Oracle Technology Network.
  4. Unzip and run the configuration script on a supported platform, and select the type of system to configure.
  5. Go through every page in OEDA and supply values for all required fields. You cannot advance to the next page if you do not supply all of the required values. You must provide naming details and DNS, NTP, and networking details.
  6. At the end of the dialogue with OEDA, configuration files are generated on the client. The files are also listed at the bottom of the `InstallationTemplate.html` file that is generated by OEDA. Depending on your engineered system and configuration, OEDA generates all or some of the following files:
     - `databasemachine.xml`
     - *`CustomerName-rackname`*`.xml`
     - *`CustomerName-rackname`*`-preconf_`*`GUID`*`.csv`
     - *`CustomerName-rackname`*`-InstallationTemplate.html`
     - *`CustomerName-rackname`*`-platinum.csv`
     - *`CustomerName-rackname`*`-checkip.sh`
     - *`CustomerName-rackname`*`.zip`
     - `pkey_`*`GUID`*`.csv` and `pkey_`*`racknamehostname_GUID`*`.csv` — if you enabled InfiniBand partitioning for your virtual environments

     The *`CustomerName-hostname`*`.zip` file contains all the generated files.

7. Review the `InstallationTemplate.html` file to check the entire configuration and verify all information was entered correctly.

- Shortly before your engineered system arrives, or is scheduled to be configured, validate the network configuration, as directed by Oracle. See Verifying the Network Configuration Prior to Configuring the Rack.
- After your engineered system arrives, the configuration files are copied to a database server, and the validation and installation is completed. See #unique_90.

Accessibility Note: If you are using Oracle Exadata Deployment Assistant with assistive technology (for accessibility) on Windows, you need to enable Java Access Bridge.

Considerations and Requirement for Using OEDA

Review this information before using Oracle Exadata Deployment Assistant (OEDA).

- ZDLRA Rack ships with the Oracle Linux operating system installed on the servers.
- The rack name is used to generate host names for all components. For example, if you specify `dbm0` for the rack name:
  - The database server host names are like `dbm0db01`
  - The storage server host names are like `dbm0cel01`
  - The InfiniBand Network Fabric switch names are like `dbm0sw-iba1`
  - The RoCE Network Fabric switch names are like `dbm0sw-rocea1`

  In you have multiple racks, then, each ZDLRA Rack name should be a unique name that identifies the machine. Oracle recommends using `dm01` for the first ZDLRA Rack, `dm02` for the second, `dm03` for the third, and so on. Up to eight ZDLRA Racks can be cabled together without additional InfiniBand switches.
- The backup method and redundancy protection levels you use determine how to size the Oracle Automatic Storage Management (Oracle ASM) disk groups that are created during installation.
  - If you configure internal backups, then the database backups are created on disk in the Fast Recovery Area in the RECO disk group for the database. The Fast Recovery Area also contains Archived Redo Logs and Flashback Log Files. The division of disk space between the DATA disk group and the RECO disk group is 40% and 60%, respectively.
  - If you configure external backups, then database backups are created on disk or tape media that is external to currently deployed Recovery Appliance. The Fast Recovery Area contains only objects such as Archived Redo Logs and Flashback Log Files.
- A valid time zone name is required when installing Recovery Appliance. Time zone values provided for Recovery Appliance and Oracle Linux comes from the Time Zone Database. You must specify a valid time zone name for each server using the `TZ` environment variable. The variable value is of the form *Area/Location*. For example, `America/New_York` is a valid `TZ` value; `EST`, `EDT`, and `UTC-4` are invalid values.

Related Topics

- iana Time Zone Database

- Maximum Availability with Oracle ASM

OEDA Browser-based User Interface

The Web-based interface for Oracle Exadata Deployment Assistant (OEDA) is available starting with the release October 2018 of OEDA. OEDA Web is available for Linux, OSX and Windows.

You can use this interface to enter configuration details for new racks and also import configuration files from the old OEDA interface. Use the **Menu** button to import and save configurations.

To use the Web-based version of OEDA, perform the following steps.

1. Download the latest OEDA release from My Oracle Support note 888828.1.

   Extract the contents of the downloaded compressed file. When you extract the contents, it creates a directory based on the operating system, such as `linux-x64`, `macosx-x64`, or `windows-i586`, to store the extracted files. This is referred to as the OEDA_HOME directory.

2. Install the software.

   Before you can use the Web-based interface, you must install and run the Web Application Server. In the created directory, locate and run the `installOedaServer` program. You do not have to be logged in as an administrator user to run this program. Use one of the following commands, where the `-p` option specifies the port number to use:

   - On Linux, Apple, and UNIX:

     ```
     ./installOedaServer.sh -p 7072
     ```

   - On Microsoft Windows:

     ```
     installOedaServer.cmd -p 7072
     ```

   > **Note:**
   >
   > The examples use port number `7072`. However, you can specify any unused port number. It is not recommended to use port numbers less than `1024`.

   The `installOedaServer` program starts the latest version of the OEDA Web Server on the local system.

3. Access the local OEDA Web Server.

   Once the OEDA Web Server is running, you can access the Web-based application by opening a browser and entering the following URL:

   `http://localhost:`*`port_number`*`/oeda`

   In the URL, *port_number* is the network port number that you specified when starting the `installOedaServer` program. For example, `7072`.

The following topics describe the OEDA Web interface:

Using the Browser-based Version of Oracle Exadata Deployment Assistant

Use the deployment assistant for initial configuration, and when adding to an existing deployment. You can import an existing configuration when adding new racks or components.

Have the following available before running the deployment assistant:

- IP addresses
- Machine names
- DNS information
- NTP information for the rack

The following options are available from the OEDA **Menu** drop-down list:

- **About**: View information about the web interface software.
- **New**: Start a new configuration.
- **Import**: Import an existing configuration file. The file format is XML.
- **Save**: Save the current configuration to an XML file.
- **Diagnostics**: Creates a ZIP file that can be used to troubleshoot and diagnosis issues with the OEDA web interface.
- **Preferences**: Specify your preferences for how the web interface behaves.
- **Validate**: Check the current configuration for completeness and correctness.
- **CheckIp**: Checks the IP addresses.
- **Exit**: Exit OEDA

The following list describes the configuration information you specify in OEDA:

1. Select the hardware for the deployment on the **Select Hardware** page. Also provide customer information on this page.
2. Enter the management network and private network information on the **Rack Networks** page.
3. Specify the user and group configuration on the **Users and Groups** page. You can choose between Default, Role Separated, or a Custom configuration.
4. Define the clusters on the **Define Clusters** page for each cluster.
5. Define the disk group layout on the **Diskgroups** page.
6. Specify the Oracle Database home information for each Oracle home on the Create **Database Home** page.
7. Enter the client and backup network information on the **Cluster Networks** page.

   If an ingest network or replication network is needed, specify their information on the **Cluster Networks** page.
8. Configure the email alerts, SNMP alerts, Auto Service Request, Oracle Config Manager, and Enterprise Manager alerting on the Alerting page.

9. Answer the configuration questionnaire and provide additional information about the deployment in the text field on the **Comments** page. Information entered in this field appears at the bottom of the `InstallationTemplate.html` file.
10. Click **Save and Download** to create the configuration files. You are prompted to choose a location in which to save a ZIP file to your local computer.
11. Copy the ZIP file to a directory on the first database server in Oracle Exadata Rack. Use a subdirectory of `/u01`, such as the `/u01/oeda` directory. These files are used by the field engineer to configure the Oracle Exadata Rack.

Hardware Selection Page

The hardware selection page identifies the number of racks that you want to deploy at the same time, and the number of storage servers within each rack. Recovery Appliance supports a flexible configuration. From the a minimum configuration X5 through X9M, you can add up to 15 storage servers; up to 11 storage servers for X4.

> **Note:**
>
> Multiple racks must be cabled together according to the *Engineered Systems Extending and Multi-Rack Cabling Guide*.

The following table describes the hardware choices.

Table 4-1 Hardware Selection Page

Hardware Selection Field	Description
Select interconnected hardware to deploy	Lists the available hardware configurations for a rack. Each rack can have 3 to 18 storage servers (cells), depending on your hardware version. Choose one or more racks. You can choose the same type of rack multiple times.
This is your deployment	Lists the hardware selected for your site.

X9M, X8M, X8, or X7 hardware: To deploy two full racks (18 storage servers) and one minimally configured rack (3 storage servers) at the same time:

1. From the **Select Engineered System** area, choose the option for **Zero Data Loss Recovery Appliance**
2. In the **Select Rack** area, use the drop-downs **ZDLRA Server** and **Capacity**.
   - X9M
   - X8M
   - X8-2
   - X7-2
   - X6-2
   - X5-2
   - X4-2

3. Select the **+ (Add selected rack)** button.

   The Rack Details page appears.

4. Fill in details about the rack being deployed.

Choose Operating System Page

Select the operating system to install on the servers in the rack.

For ZDLRA racks, your choices are:

- **All Linux Physical**: All compute nodes installed with Oracle Linux without Oracle VMs.

Rack Networks Page

Provide the information necessary to configure the administration and private networks on the rack.

Each compute node and storage server has two network interfaces for administration. One network provides management access to the operating system through the Ethernet interface, and the other network provides access to the Integrated Lights Out Manager (ILOM) through the ILOM Ethernet interface. Default administration network host names are usually changed to meet corporate naming standards or compatibility with existing systems.

In the upper right corner of the Admin Network section is a line that displays the Total IP Addresses that will be configured for the rack. This value is based on the hardware you selected on the Hardware page.

You provide the following information on this page:

Admin Network

Supply values for the following fields.

- **Default gateway for database servers**: A check box that indicates that the gateway IP address is the default gateway for the database servers.
- **Gateway**: Sets the IP address for the gateway. This field cannot be empty.
- **Start IP Address**: Assigns the starting IP address for the administration network of the servers and switches. This field cannot be empty.

  Oracle Exadata Deployment Assistant (OEDA) generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified during the configuration process.

> **Note:**
>
> These addresses are also assigned to the ILOM ports on the database and storage servers

- **Subnet Mask**: Sets the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from `255.0.0.0` to `255.255.255.252`
- **Domain Name**: Sets the network domain name. Use lowercase letters for the name. This field cannot be empty.

  The domain name specified on this page is used as the domain for all networks. If different domains are used for the client and backup networks, then specify the administration network on this page. The client and backup networks can be changed on the Cluster Networks page.
- **Separate ILOM Network**: A check box that indicates that ILOM uses a network that is different from the administration network. If you check this option, you must specify the following additional information for the ILOM network:
  - **Gateway**
  - **Start IP Address**
  - **Subnet Mask**
  - **Domain Name**

Advanced

If you click the **Advanced** button in the upper right corner of this page, you can select the following options:

- **Enable Vlan** — Allows you to create VLANs on database servers and storage servers for the admin network, ILOM, client and the backup access network. If you select this option, then the following additional fields are visible on the Admin Network page:
  - **Vlan** — Specifies the VLAN identifier. The value must be a positive integer.

> **Note:**
>
> * If your system will use more than 10 VIP addresses in the cluster and you have VLAN configured for the Oracle Clusterware client network, then you must use 3 digit VLAN ids.
> * Do not use 4 digit VLAN ids because the VLAN name can exceed the 15 character operating system interface name limit.
> * If VLAN tagging is to be used, then it must be enabled by the customer after the initial deployment.
>
> See the My Oracle Support note Enabling VLAN Tagging in Recovery Appliance (Doc ID 2047411.1) for details.

Modify Masks

If you click the button **Modify Masks**, it opens a pop-up window where you can enter the following information:

- Compute Name Mask — There is a preview field at the bottom of this window that shows how the generated administration names will appear based on the values you specify for the following fields:

> **Note:**
>
> In the following fields, the name prefix can contain alphanumeric characters. The percent signs (%) in the names are replaced by the numbers entered into the Start ID field. At least one percent sign is required in each mask field. If you use more than one percent sign in a mask field, the symbols must be contiguous. For example, you can use `test%%-ilom`, but not `test%new%-ilom`.

  - **Compute Name Mask**: The string used to generate administration names for each database server in this rack. The default string is {*Rack Prefix*}adm%%, which results in names such as dbm0adm01, dbm0adm02, and so on. This field cannot be empty.
  - **Compute Name Start Id**: The starting value to use when generating the database server administration names. The default value is 1. This field cannot be empty.

- Compute ILOM Mask — There is a preview field at the bottom of this section that shows how the generated administration names will appear based on the values you specify for the following fields:
  - **Compute ILOM Mask**: The string used to generate administration names for each database server ILOM in this rack. The default string is {*Compute Name Mask*}%-ilom, which results in names such as dbm0adm01-ilom, dbm0adm02-ilom, and so on. This field cannot be empty.
  - **ILOM Name Start Id**: The starting value to use when generating the ILOM administration names. The default value is the same value as Compute Name Start Id. This field cannot be empty.

> **Note:**
>
> Oracle recommends using the same starting IDs for the servers and server ILOMs.

- **Storage Name Mask** — There is a preview field at the bottom of this section that shows how the generated administration names will appear based on the values you specify for the following fields:
  - **Storage Name Mask**: The string used to generate administration names for each storage server. The default string is {*Rack Prefix*}celadm%%, which results in names such as dbm0celadm01, dbm0celadm02, and so on. This field cannot be empty.
  - **Storage Name Start Id**: The starting value to use when generating the storage server network names. The default value is 1. This field cannot be empty.
- **Storage ILOM Mask** — There is a preview field at the bottom of this section that shows how the generated server names will appear based on the values you specify for the following fields:
  - **Storage ILOM Mask**: The string used to generate administration names for each storage server ILOM in this rack. The default string is {*Storage Name*

Mask}%-ilom, which results in names such as dbm0celadm01-ilom, dbm0celadm02-ilom, and so on. This field cannot be empty.

  - **Storage Name Start Id**: The starting value to use when generating the ILOM network names. The default value is the same value as Storage Name Start Id. This field cannot be empty.

> **Note:**
>
> Oracle recommends using the same starting IDs for the servers and server ILOMs.

- Switches and PDUs
  - **Cisco Switch** — The string used to generate the administration name for the Cisco Management Switch in this rack. The default value is {*Rack Prefix*}sw-adm%, which results in names such as dbm0sw-adm1.
  - **RDMA Network Fabric Spine Switch** — The string used to generate the administration names for the spine switch, if your rack is configured with one. The default value is {*Rack Prefix*}sw-ibs%, which results in names such as dbm0sw-ibs1.
  - **RDMA Network Fabric Leaf Switch** — The string used to generate the administration names for the lower RDMA Network Fabric Leaf switch in this rack. The default value is {*Rack Prefix*}sw-*type*a%, which results in names such as dbm0sw-iba1 or dbm0sw-rocea0.
  - **RDMA Network Fabric Leaf Switch** — The string used to generate the administration names for the upper RDMA Network Fabric Leaf switch in this rack. The default value is {*Rack Prefix*}sw-*type*b%, which results in names such as dbm0sw-ibb1 or dbm0sw-roceb0.
  - **PDU-A** — The string used to generate the administration names for Power Distribution Unit (PDU) A in this rack. The default value is {*Rack Prefix*}sw-pdua%, which results in names such as dbm0sw-pdua1.
  - **PDU-B** — The string used to generate the administration names for Power Distribution Unit (PDU) B in this rack. The default value is {*Rack Prefix*}sw-pdub%, which results in names such as dbm0sw-pdub1.

After you have entered the custom network name masks, click the **Save** button at the bottom of the page. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the page.

Admin network Review and Edit

After you save the input provided for the admin network, the Admin network Review and Edit section appears above the Private Network section. Here you can view all the IP addresses and administration names that will be configured for your rack, and modify them if needed. If you make any modifications to the administration name or their IP addresses, click the **Update** button when you have finished to save your changes.

If you are connected to the same network on which the rack will be deployed, you can check the validity of the network configuration data. Click the **Lookup IP** button at the bottom of this section. OEDA will display any invalid IP addresses.

Private Network

Provide the information necessary to configure the private network on the rack.

In the upper right corner of the Private Network section is a line that displays the Total IP Addresses that will be configured for the rack. This value is based on the hardware you selected on the Hardware page.

- **Enable active bonding** — Specifies whether active bonding should be enabled for only the private network on compute nodes. By default this option is selected.
- **Start IP Address**: Assigns the starting IP address for the private network of the servers. This field cannot be empty.

  OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified during the configuration process.
- **Subnet Mask**: Sets the IP address mask for the subnet. You can choose a subnet mask from the drop-down list. The list contains supported values from `255.255.240.0` to `255.255.255.252`
- **Domain Name**: Sets the network domain name. The default value is the domain name value entered for the Admin Network. Use lowercase letters for the name. This field cannot be empty.

If you click the **Modify Masks** button, it opens a popup window where you can modify the generated names for the database and storage server private names. There are preview fields at the bottom of this window that shows how the generated private names will appear based on the values in the name mask fields. The default strings used to generate the private names are:

- Compute private name mask: {*Rack prefix*}`adm%%-priv`, which results in names such as `dbm0adm01-priv`, `dbm0adm02-priv`, and so on
- Cell private name mask: {*Rack Prefix*}`celadm%%-priv`, which results in names such as `dbm0celadm01-priv`, `dbm0celadm02-priv`, and so on.

After you have entered the custom private network name masks, click the **Save** button at the bottom of the page. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the page.

After you have entered the private network configuration information, click the **Save** button at the bottom of the page.

Private network Review and Edit

After you save the input provided for the private network, the Private network Review and Edit section appears above the User(s) and Groups section. Here you can view all the IP addresses and private host names that will be configured for your rack, and modify them if needed. If you make any modifications to the private host names or their IP addresses, click the **Update** button at the bottom of this section when you have finished to save your changes.

Users and Groups

Provide information about the operating system users to use when installing Oracle software. The user information you provide is saved in a collection, which can be referenced later in the configuration process.

The first selection is the type of users you want to define. Your options are:

- **Default**: Use the default user configuration. This option creates a single operating system user to own the Oracle Grid Infrastructure and Oracle Database software.

If you plan to have multiple installations of Oracle software on your rack that require different operating system users, then, after you have entered the user information, you can click the **Add group (+)** button. The information you entered is saved with the label `UserCollection-1_id`. A new section appears on the page where you can enter information for an additional group of software installation users with distinct names and permissions. The new user information collection is automatically assigned the label `UserCollection-2_id`.

If you want to remove a user information collection, click the **Delete Group (-)** button next to the ID of the group you want to remove.

After you have entered all the user information, click the **Save** button at the lower right side of the page to continue.

Default User Information

For ZDLRA racks, the **Default** option is selected automatically, and the only values you can modify are the **User Id** and **Group Id**.

- **User Name**: The operating system user name for the database software owner. This field cannot be empty. The default value is `oracle`.
- **User Id**: The operating system identification number for this user. This field cannot be empty. The default value is `1001`.
- **Home Directory**: The user home directory. This field cannot be empty. The default value is `/home/oracle`.
- **User Type**: This field is not available if you select the Default option. The named user will be created as the Grid and Database software owner.
- **Group Name**: The name of the operating system group that will be the DBA group. All the operating system privileges necessary for managing the Oracle Database software are granted to this group. This field cannot be empty. The default value is `dba`.
- **Group Id**: The operating system identification number for this group. This field cannot be empty. The default value is `1001`.
- **Group Type**: This field is not available if you select the Default option. The named group will be created as a `DBA GROUP` group.
- **Group Name**: The name of the operating system group that will be the installation owner group. All the operating system privileges necessary for installing the Oracle Grid Infrastructure and Oracle Database software are granted to this group. This field cannot be empty. The default value is `oinstall`.
- **Group Id**: The operating system identification number for this group. This field cannot be empty. The default value is `1002`.
- **Group Type**: This field is not available if you select the Default option. The named group will be created as an `OINSTALL` group.

Save

After you have entered all the information for the racks being configured, click the **Apply** button. Be sure to go to the **Menu** and select **Save**.

Define Clusters

More than one cluster can be defined for a hardware racks. Larger environments support multiple clusters. Each cluster requires a minimum of two physical or virtual compute nodes, and three storage cells. You cannot mix physical and virtual compute nodes in the same cluster.

Provide the information for the first cluster and assign nodes to the cluster using the **Add Node (>)** button to move selected nodes or the **Add All Nodes (>>)** button to assign all available nodes to the cluster. If you want to add a second cluster, then click the **Add Cluster (+)** button next to the cluster name tab.

Provide the following information for each cluster in the rack:

- **Cluster Name**: Is the cluster name, which must be unique across your enterprise. The cluster name is case-insensitive, must be at least one character long and no more than 15 characters in length, must be alphanumeric, cannot begin with a numeral, and may contain hyphens (-). Underscore characters (_) are not allowed. This field cannot be empty.
- **Inventory Location**: Is the Oracle inventory (`oraInventory`) directory path. The default value is `/u01/app/oraInventory`.
- **Grid Home Owner**: Use the drop-down list to select the collection of user information previously created on the Users page. By default, the first user collection is selected.
- **Grid Infrastructure Home Location**: Is the Oracle Grid Infrastructure directory path, which typically includes the software release number. The default value is `/u01/app/`*`release_number`*`/grid`, for example, `/u01/app/18.2.0.0/grid`. This value is automatically set depending on the value selected for Cluster Version. You can overwrite the default value.
- **Cluster Version**: After you select the nodes to assign to this cluster, you will be able to select the cluster version from the drop-down list. The list contains all supported versions of the Oracle Grid Infrastructure. Choose any of the supported values.
- **Base Directory**: Is the base directory for the Oracle Grid Infrastructure software. The default value is `/u01/app/grid`.

After you assign nodes to the cluster, you can modify the default values for each cluster, if needed. The following customer information fields appear below the Available Machines list:

- **Customer Name**: Is the customer name. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **Application**: Is the application name or description. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **Region**: Is the geographical location of the rack. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **Time zone**: Is the time zone in which the rack is located. The default value is the value you provided earlier on the Hardware page in the Customer Details section.

- **DNS Servers**: Are the IP addresses for the DNS servers. The default values are the values you provided earlier on the Hardware page in the Customer Details section.
- **NTP Servers**: Are the IP addresses for the NTP servers. The default values are the values you provided earlier on the Hardware page in the Customer Details section.

Diskgroups

Define the Oracle ASM disk groups and their sizes for each cluster in the rack.

If you specified more than one cluster, use the tabs at the top of this page to select each cluster and provide the disk group information.

For ZDLRA racks, the following information is used:

- Diskgroup name: Automatically set to CATALOG and DELTA
- Diskgroup Type: Automatically set to CATALOG and DELTA
- Redundancy: The CATALOG disk group is configured to use HIGH redundancy. For the DELTA disk group, you can choose either HIGH or NORMAL (default). HIGH reduces the usable space.
- Size: The CATALOG disk group has a size of Default, and the DELTA disk group has a size of 95%.
- Size Type: The disk group sizes are specified as percent.

After you have entered the values for the disk groups and clicked Save, the following fields are then populated with the estimated values:

- CATALOG disk group — The usable space and raw size are set to default.
- DELTA disk group — The values depends on the hardware selected.

Database Home

Provide information for configuring the Oracle Database software installation.

For ZDLRA hardware, you should accept all default values for Oracle Database software installation:

- **Select Cluster**: The cluster on which to deploy the Oracle Database software. This is a drop-down list that is populated with the cluster names specified on the Clusters page. By default, the first cluster is selected.
- **Database Home Name**: A name given to the Oracle Database software installation. This name identifies the program group associated with a particular home and the Oracle services installed on this home. The Oracle home name must be between 1 to 127 characters long, and can include only alphanumeric characters and underscores. This field cannot be empty. The default value is `DbHome_1`.
- **Database Home Version**: The version of the Oracle Database software to install in the cluster. Choose the value from the drop-down list. The most recent version is selected by default.
- **Database Home Location**: A directory path where the Oracle Database software will be installed. This value is set automatically based on the value specified for the Database Home Version. For example, if you select version `12.1.0.2 DBBP13`, then the Database Home Location is set to `/u01/app/oracle/product/12.1.0.2/dbhome_1`. After

specifying the Database Home Version, you can type over the default value to use a different location, if needed. This field cannot be empty.

- **Base Directory**: The Oracle Base directory location. The Oracle Database binary home is located underneath the Oracle base path. You can use the same Oracle base directory for multiple installations of Oracle software. If different operating system users install Oracle software on the same system, then you must create a separate Oracle base directory for each installation owner. This field cannot be empty. The default value is `/u01/app/oracle`.

When you have finished configuring the Oracle home information, click **Apply** and then **Save** from the **Menu**.

Database

Provide information about the databases to create in the specified database home directories.

- **Database Name**: Is the name of the database you want to create in the database home. The database name (`DB_UNIQUE_NAME`) portion is a string of no more than 30 characters that can contain alphanumeric, underscore (_), dollar ($), and pound (#) characters, but must begin with an alphabetic character. No other special characters are permitted in a database name. The `DB_NAME` parameter for a database is set to the first 8 characters of the database name. The domain portion of the global database name (`DB_DOMAIN`) uses the value you provided for the domain when configuring the networks. The values for `DB_UNIQUE_NAME.DB_DOMAIN` in its entirety must be unique within the enterprise. This field cannot be empty.

  `DB_UNIQUE_NAME` is filled in from previous domain name selection and does not change. The database name field **must** be unique across the enterprise Recovery Appliances. The resulting **Database Name** pulls previously entered domain name and appends to user selected `DB_UNIQUE_NAME`.

- **Database Home**: Is the name of the Database Home in which you want to create the database. This value is automatically populated with the first home specified on the Database Home page.

When you have finished configuring the database configuration information, click **Save** from the **Menu**.

Cluster Networks

For ZDLRA racks, configure the Ingest and Replication networks.

Applications access the database servers or compute nodes through the client Ethernet network using Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) addresses. Oracle databases running in Oracle VM guests on Oracle Exadata Database Machine are accessed through the client Ethernet network as defined in this section. Client network configuration in both the management domain (dom0) and user domains (domUs) is done automatically when the OEDA installation tool creates the first user domain during initial deployment.

In the upper right corner of the Cluster Networks page is a line that displays the Total IP Addresses that will be configured for the rack. This value is based on the hardware you selected on the Hardware page. At the top of the section are tabs, each designating one of the clusters you specified on the Clusters page.

Advanced

At the very top of the Cluster Networks section, before any network configuration details, there is a button for **Advanced** configuration. If you click this button, the Advanced options window appears. Depending on the system configuration entries that you have made so far, you may see the following options:

- **Enable Vlan**: Select this option if there is a need for virtual deployments on the rack to access additional virtual local area networks (VLANs) on the network, such as enabling network isolation across user domains.

  If you select this option, then in the client or ingest network configuration details, an additional field appears:

  – Vlan: Specify the VLAN identifier. The VLAN ID is a numeric value between 1 and 4095. You also need to select the `Bonded` option to enable VLAN. If the Backup or Replication network uses VLAN, then the Client or Ingest network must also be configured to use VLAN. If the VLAN ID field is empty, then VLAN is not enabled.

> **Note:**
>
> For all networks, you are responsible for configuring VLAN on the switch.

Admin Network

For some Exadata rack configurations, at the top of the Cluster Networks section, on the right side, is a check box **Enable Admin network**. You can use this to optionally configure an Admin network inside the virtual compute servers. If you select this option, you must specify the network configuration by using the following:

- **Default gateway for database servers**: Select this option if the gateway IP address is the default gateway for all database servers
- **Defines Domain hostname for database servers**: Select this option if you want the name used for the Admin network on a virtual compute server to be used as the host name for the virtual compute server. If you do not select this option, then the virtual compute server uses the network name assigned to the client network as the host name.
- **Gateway**: Sets the IP address for the gateway for the client network. This field cannot be empty.
- **Start IP Address**: Specify the first IP address that will be used for client access to the database servers. This value is filled in automatically, based on the value supplied for the gateway. OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified using the **Modify Masks** button. This field cannot be empty.
- **Subnet Mask**: Specify the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from `255.0.0.0` to `255.255.255.252`.
- **Domain Name**: Specify the network domain name for client connections. Use lowercase letters for the name. This field cannot be empty. For Exadata Cloud configurations, this value is provided for you.

After you have specified all the configuration details for the Admin network, click the **Apply** button at the bottom of the section, or click **Modify Masks** to customize the network configuration.

Modify Masks

If you click the **Modify Masks** button, it opens a pop-up window where you can modify the generated names for the client or ingest network. There are preview fields at the bottom of each section that show how the generated network names will appear based on the values in the name mask fields. The default strings used to generate the network names are:

- Client Access Details:
  - **Name**:
    - ZDLRA racks: Enter {*Rack prefix*}ingest%%, which results in names such as {*Rack prefix*}ingest01, {*Rack prefix*}ingest02, and so on
  - **Start id**: Specify the number to start with when generating sequential names.
- Vip details:
  - **Name**:
    - ZDLRA racks: Enter {*Rack prefix*}ingest%%-vip, which results in names such as {*Rack prefix*}ingest01-vip, {*Rack prefix*}ingest02-vip, and so on
  - **Start id**: Specify the number to start with when generating sequential names.
- Scan details: The configuration details for the single client access network (SCAN) listener
  - **Name**: Enter the name of the listener. The default value is {*Rack prefix*}-scan1.
  - **Port**: Specify the network port which the SCAN listener monitors for incoming connection requests. The default value is 1521.

After you have completed the modifications to the client or ingest network configuration, click the **Apply** button at the bottom of the Client Network Masks window. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the Client Network Masks window.

Client Network or Ingest Network

Provide the following information to configure either the Client networks for Exadata or the Ingest networks for ZDLRA racks.

For Exadata racks and ZDLRA racks:

- **Default gateway for database servers**: Select this checkbox if the gateway IP address is the default gateway for all database or compute servers.
- **Default hostname for database servers**: Select this checkbox if the network names define the host names for the database or compute servers.
- **Network Media**: Select the network cable type, either Optical or Copper.
- **Bonded**: For X7 or later racks and Exadata Cloud, only bonded networks are supported.

For X6 and earlier racks of Exadata and ZDLRA, or Exadata Cloud, select this option if using a bonded network.

- **LACP**: Check this option to enable the Link Aggregation Control Protocol (LACP) on the network.
- **Gateway**: Sets the IP address for the gateway for the network. This field cannot be empty.
- **Start IP Address**: Set the first IP address that will be used for access to the database servers on this network. This value is filled in automatically, based on the value supplied for the gateway. OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified using the **Modify Masks** button. This field cannot be empty.
- **Subnet Mask**: Set the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from `255.0.0.0` to `255.255.255.252`.
- **Domain Name**: Set the network domain name for client or ingest connections. Use lowercase letters for the name. This field cannot be empty. For Exadata Cloud configurations, this value is provided for you.

After you have specified all the configuration details for the client or ingest network configuration, click the **Apply** button at the bottom of the section.

Client network or Ingest network Review and Edit

After you save the input provided for the client or ingest network, the Client network Review and Edit or Ingest network Review and Edit section appears. Here you can view all the host names and IP addresses that will be configured for client access to the database servers in the rack. You can modify any of the settings if needed. If you make any modifications to the administration name or their IP addresses, click the **Apply** button when you have finished to save your changes.

If you are connected to the same network on which the rack will be deployed, you can check the validity of the network configuration data. Click the **Lookup IP Address** button at the bottom of this section. OEDA will display any invalid IP addresses.

Backup Network or Replication Network

In the upper right corner of the Backup Network section for Exadata is the option **Enable Backup network**. For ZDLRA racks, this option is labeled **Enable replication network**. If you select this option, it enables the configuration of a network dedicated to backups or replication.

In the upper right corner of the Backup Network or Replication Network section is a line that displays the Total IP Addresses that will be configured for each cluster. At the top of the section are tabs, each designating one of the clusters you specified on the Clusters page.

Provide the following information to configure the backup network or replication network for each cluster:

- **Shared client network port**: (Not available for Exadata Cloud) Select this option if the client and backup networks use the same physical network cables. This option requires that both networks have VLAN enabled and share the same network interface cards (NIC). If this option is not selected, then you must use one set of cables for the client network and a separate set of cables for the backup network.

- **Network Media**: (For Exadata X6 racks and earlier and ZDLRA racks) Select the network cable type, either Optical or Copper.
- **Bonded**: (Not available on X7 or later racks or Exadata Cloud) Select this option to use network bonding. For X7 or later racks for Exadata and ZDLRA, and Exadata Cloud, only bonded networks are supported.
- **LACP**: Check this option to enable the Link Aggregation Control Protocol (LACP) on the network.
- **Gateway**: Set the IP address for the gateway for the network. This field cannot be empty.
- **Start IP Address**: Set the first IP address that will be used for the network. This value is filled in automatically, based on the value supplied for the gateway. OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified using the **Modify Masks** button. This field cannot be empty.
- **Subnet Mask**: Set the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from 255.0.0.0 to 255.255.255.252.
- **Domain Name**: (Not available for Exadata Cloud) Set the network domain name for network connections. Use lowercase letters for the name. This field cannot be empty.

After you have specified all the configuration details for the network configuration, click the **Apply** button at the bottom of the section.

If you click the **Modify Masks** button, it opens a pop-up window where you can modify the generated names for the network. There is a preview field at the bottom of the window that shows how the generated network names will appear based on the values in the name mask fields. The default strings used to generate the network names are:

- **Client Access Details**:
  - **Name**:
    - ZDLRA racks: Enter {*Rack prefix*}`ingest%%`, which results in names such as {*Rack prefix*}`ingest01`, {*Rack prefix*}`ingest02`, and so on
  - **Start id**: Set the number to start with when generating sequential names.
- **Vip details**:
  - **Name**:
    - ZDLRA racks: Enter {*Rack prefix*}`ingest%%-vip`, which results in names such as {*Rack prefix*}`ingest01-vip`, {*Rack prefix*}`ingest02-vip`, and so on
  - **Start id**: Set the number to start with when generating sequential names.
- **Scan details**: Are the configuration details for the single client access network (SCAN) listener.
  - **Name**: Specify the name of the listener. The default value is {*Rack prefix*}`-scan1`.
  - **Port**: Specify the network port which the SCAN listener monitors for incoming connection requests. The default value is 1521.

After you have completed the modifications to the backup or replication network configuration, click the **Apply** button at the bottom of the Backup Network Masks or Replication Network Masks page. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the window. The buttons might be out of view are require scrolling down in the window to see.

After you have specified all the configuration details for the backup or replication network configuration, click the **Apply** button at the bottom of the Backup Network or Replication Network section.

Backup or Replication network Review and Edit

If you enabled the Backup network or the Replication network option, then this section appears after you save the input provided for the network. Here you can view all the host names and IP addresses that will be configured for the network. You can modify any of the settings if needed. If you make any modifications, then click the **Update** button when you have finished to save your changes.

If you are connected to the same network on which the rack will be deployed, you can check the validity of the network configuration data. Click the **Lookup IP Address** button at the bottom of this section. OEDA will display any invalid IP addresses.

Alerting

In this section you configure alerting

You use the following methods for receiving or handling alerts:

- Email with Simple Mail Transfer Protocol (SMTP)
- Simple Network Management Protocol (SNMP)
- Oracle Auto Service Request (ASR)
- Oracle Configuration Manager
- Oracle Enterprise Manager

For ZDLRA, you must configure at least one of the following alerting options:

- Oracle ASR
- SMTP or SNMP alerting
- Oracle Configuration Manager
- Oracle Platinum Services

The alerting options include:

- **Enable Email Alerting**: Enables alerts to be delivered automatically through email.
  - **Name**: Sets the SMTP email user name that appears in the alert notifications, for example, `Oracle Exadata Database Machine`. This field cannot be empty.
  - **Email Address**: Sets the SMTP email address that sends alert notifications, for example, `dm0db01@example.com`. This field cannot be empty.
  - **SMTP Server**: Sets the SMTP email server used to send alert notifications, for example, `mail.example.com`. This field cannot be empty.
  - **Port**: Sets the SMTP email server port used to send alert notifications, for example, `25` or `465`. This field cannot be empty.

  - **Send Email Alerts to**: Sets the email address for the recipients of the alerts. You can click the **Add Email (+)** button to add additional recipients. You can use the **Remove Email (-)** button to remove an email address that you added.

- **Enable SNMP Alerting**: Enables using SNMP to deliver alerts.
  - **SNMP Server**: Sets the SNMP server host name, for example, `snmp.example.com`. This field cannot be empty.

    > **Note:**
    >
    > Additional SNMP targets may be defined after installation of the Exadata rack. Refer to the `ALTER CELL` CellCLI command or the `ALTER DBSERVER` DBMCLI command.

  - **Port**: Sets the SNMP email server port used to send alert notifications. The default value is `162`. This field cannot be empty.
  - **Community**: Sets the SNMP server community. The default value is `public`. This field cannot be empty.
- **Enable ASR Alerting**: Enables Oracle Auto Service Request (ASR). Oracle ASR is designed to automatically open service requests when specific rack hardware faults occur.

  Oracle recommends that you install Oracle ASR Manager on an external, standalone server. The server must be installed have connectivity to the Exadata servers or Recovery Appliance nodes and an outbound Internet connection using HTTPS or an HTTPS proxy. This server must run an Oracle Solaris or Linux operating system.
  - **ASR Manager Request**: The host name of the Oracle ASR Manager server. This field cannot be empty.
  - **HTTP Proxy Host**: If you want to enable an HTTP proxy to upload to Oracle ASR, then specify the proxy server host name.
  - **Contact Name**: The name of the technical contact person for the rack. This field cannot be empty.
  - **Contact Email**: The e-mail address of the technical contact for the rack. This field cannot be empty.
  - **HTTP Proxy User**: If you want to enable an HTTP proxy to upload to Oracle ASR, then specify the proxy server user name.
  - **Port**: If you want to enable an HTTP proxy to upload to Oracle ASR, then specify the HTTP proxy server port number.
  - **My Oracle Support Account Name**: The email address of the My Oracle Support account. This field cannot be empty.
- **Enable OCM Alerting**: Enables Oracle Configuration Manager to collect configuration information.

  Oracle Configuration Manager is used to collect configuration information and upload it to an Oracle repository, for access when needed by Oracle Support Services.
  - **MOS Email Address**: The My Oracle Support account (email address) to receive updates from My Oracle Support. This field cannot be empty.

  - **HTTP Proxy Host**: If you use an HTTP proxy to upload configuration information to the Oracle repository, then specify the proxy server host name. The proxy server may be between the following:
    * Database (compute) servers and Oracle (applicable when you do not have Oracle Support Hub)
    * Database (compute) servers and Oracle Support Hub (applicable when you have only Oracle Support Hub)
    * Oracle Support Hub and Oracle

    If passwords are needed, then provide them during installation.
  - **Support Hub Hostname**: The Support Hub server host name. Oracle Support Hub allows Oracle Configuration Manager instances to connect to a single internal port (the Oracle Support Hub), and upload configuration data. This eliminates the need for each individual Oracle Configuration Manager instance in the database servers to access the Internet. This field cannot be empty.
  - **HTTP Proxy User**: If you want to enable an HTTP proxy to upload configuration information to the Oracle repository, then specify the HTTP proxy server user name.
  - **Hub User Name**: The Oracle Support Hub operating system user name. This field cannot be empty.
  - **Port**: If you want to enable an HTTP proxy to upload configuration information to the Oracle repository, then specify the HTTP proxy server port number. The default value is `80`.

- **Enable EM Alerting**: Enables alerts in Oracle Enterprise Manager.

  Oracle Enterprise Manager Cloud Control or Oracle Enterprise Manager Grid Control is used to monitor the rack. When Oracle Enterprise Manager agents are installed on the rack, they can send information to the Oracle Management Server. If you do not have an existing Oracle Management Server, then install a new Oracle Management Server on a separate host.

  - **EM Home Base Location**: The base directory for the Oracle Enterprise Manager agents installed on the database or compute servers. The default value is `/u01/app/oracle/product/EMbase`. This field cannot be empty.
  - **OMS HTTPS Upload Port**: The port the Oracle Enterprise Manager agent should use to upload information to the Oracle Management Server. The default value is `4900`. This field cannot be empty.
  - **OMS Host Name**: The name of the server or host where the Oracle Management Server repository is installed. This field cannot be empty.

Tape Library Connectivity

On ZDLRA racks only, you can configure network connection information for tape backups.

You can configure the tape library for ZDLRA only if you use Oracle Secure Backup as the media manager for tape backups. You can manually configure other media managers as clients only.

- **Use Tape**: Select this option if you plan to use Oracle Secure Backup as the media manager for tape backups. Then you can enter the configuration details.
- **Vip IP Address**: The IP address of the VIP.

- **Vip Name**: The alphanumeric name of the VIP. The default value is `osbadminvip.{domain_name}`.

  If you are connected to the same network on which the rack will be deployed, you can check the validity of the tape library address information. Click the **Lookup IP** button at the bottom of this section. OEDA will display any invalid IP addresses.
- **Number of tape libraries**: The number of tape libraries available to the ZDLRA rack. The default is 1. All tape libraries must be hosted in the same domain. If you select a value greater than 1, a new tab appears for each tape library. You can enter the configuration details for each tape library on the separate tabs.
- **ACLS ACS API**: Use the Automated Cartridge System Library Software (ACSLS) API for the tape library. Fill out the following additional fields:
  - **ACSLS Server**: The name of the server running ACSLS
  - **ACS ID**: The ACS identification number
- **ACSLS FC-SCSI**: Use the Automated Cartridge System Library Software (ACSLS) FC-SCSI for the tape library. Fill out the following additional fields:
  - **ACSLS Server**: The name of the server running ACSLS
  - **Logical Library ID**: The logical library identification number
- **SCSI**: Use SCSI to communicate with the tape library. The Simple SAN tape library has no configuration fields. This is the default option.

Comments

Enter any additional information about the deployment in the text field.

For Oracle Exadata Rack, this section contains a list of questions about the configuration to help ensure a smooth installation and configuration process for your rack. You can type in the answers in the text field.

For all deployments, you can provide additional comments about your environment or the deployment in this field.

When you have finished providing the information, click **Save and Download**. You are prompted to save a ZIP file containing the deployment and configuration files. The main configuration files are:

> **Note:**
>
> Do not edit or modify any of the generated files.

- *Customer_name-rack_prefix*`.xml`: This is the main configuration file used by the installation process.
- *Customer_name-rack_prefix*`-checkip.sh`: A script you run on the installed rack to verify all the IP addresses you specified in OEDA are available and ready for deployment.
- *Customer_name-rack_prefix*`-InstallationTemplate.html`: This file is an HTML report of all the data captured in the web user interface. The questions and answers from the Comments section of OEDA appear at the bottom of this file.

- *Customer_name-rack_prefix*-platinum.csv: This file is used when first imaging of the nodes in the rack by Oracle Advanced Customer Support.
- *Customer_name-rack_prefix*-preconf_*pkey*.csv: If you configured InfiniBand partitioning, this file is created. This file is used when PKEY security is enabled in the interconnect.

Part II
Installation and Configuration

Part II provides information about installation and configuration of Recovery Appliance after it has arrived on site. This part contains the following chapters:

5

Installing the Recovery Appliance at the Site

This chapter describes how to move and install the hardware, from unpacking the Recovery Appliance to powering it on. This chapter contains these sections:

Reviewing Safety Guidelines

Before ZDLRA Rack arrives, the following safety precautions should be reviewed to ensure the site is safe, as well as ready for delivery.

Failing to observe these precautions can result in injury, equipment damage, or malfunction.

- Do not block ventilation openings.
- Do not install ZDLRA Rack in a location that is exposed to direct sunlight or near a device that may become hot.
- Do not install ZDLRA Rack in a location that is exposed to excessive dust, corrosive gases, or air with high salt concentrations.
- Do not install ZDLRA Rack in a location that is exposed to frequent vibrations. Install ZDLRA Rack on a flat, level surface.
- Use a power outlet that uses proper grounding. When using shared grounding, the grounding resistance must not be greater than 10 ohms. Ensure that your facility administrator or a qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.
- Be sure that each grounding wire used for ZDLRA Rack is used exclusively for ZDLRA Rack. Also be sure to observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.
- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.
- Do not place anything on top of ZDLRA Rack or perform any work directly above the rack.

- Do not let the room temperature rise sharply, especially in winter. Sudden temperature changes can cause condensation to form inside ZDLRA Rack. Allow for a sufficient warm-up period prior to server operation.
- Do not install ZDLRA Rack near a photocopy machine, air conditioner, welding machine, or any other equipment that generates loud, electronic noises.
- Avoid static electricity at the installation location. Static electricity transferred to ZDLRA Rack can cause malfunctions. Static electricity is often generated on carpets.
- Confirm the supply voltage and frequency match the electrical ratings indicated for ZDLRA Rack.
- Do not insert anything into any ZDLRA Rack opening, unless doing so is part of a documented procedure. ZDLRA Rack contains high-voltage parts. If a metal object or other electrically-conductive object enters an opening in ZDLRA Rack, then it could cause a short circuit. This could result in personal injury, fire, electric shock, and equipment damage.
- When using single phase power distribution units (PDUs), note the following:
  - PDU A input 0 and PDU B input 2 must be on the same phase.
  - PDU A input 1 and PDU B input 1 must be on the same phase.
  - PDU A input 2 and PDU B input 0 must be on the same phase.

  The inputs are labeled where they come out of the PDU. Connecting cables as described ensures the phases are balanced on both sides, A and B, in case of a failover.

See Also:

- *Important Safety Information for Sun Hardware Systems* (816-7190) included with the rack
- *Oracle Engineered System Safety and Compliance Guide, Compliance Model No.: ESY27*

Unpacking ZDLRA Rack

The unpacking location should be determined during the site planning process. The following image shows the ZDLRA Rack shipping crate.

Figure 5-1 ZDLRA Rack in Shipping Crate

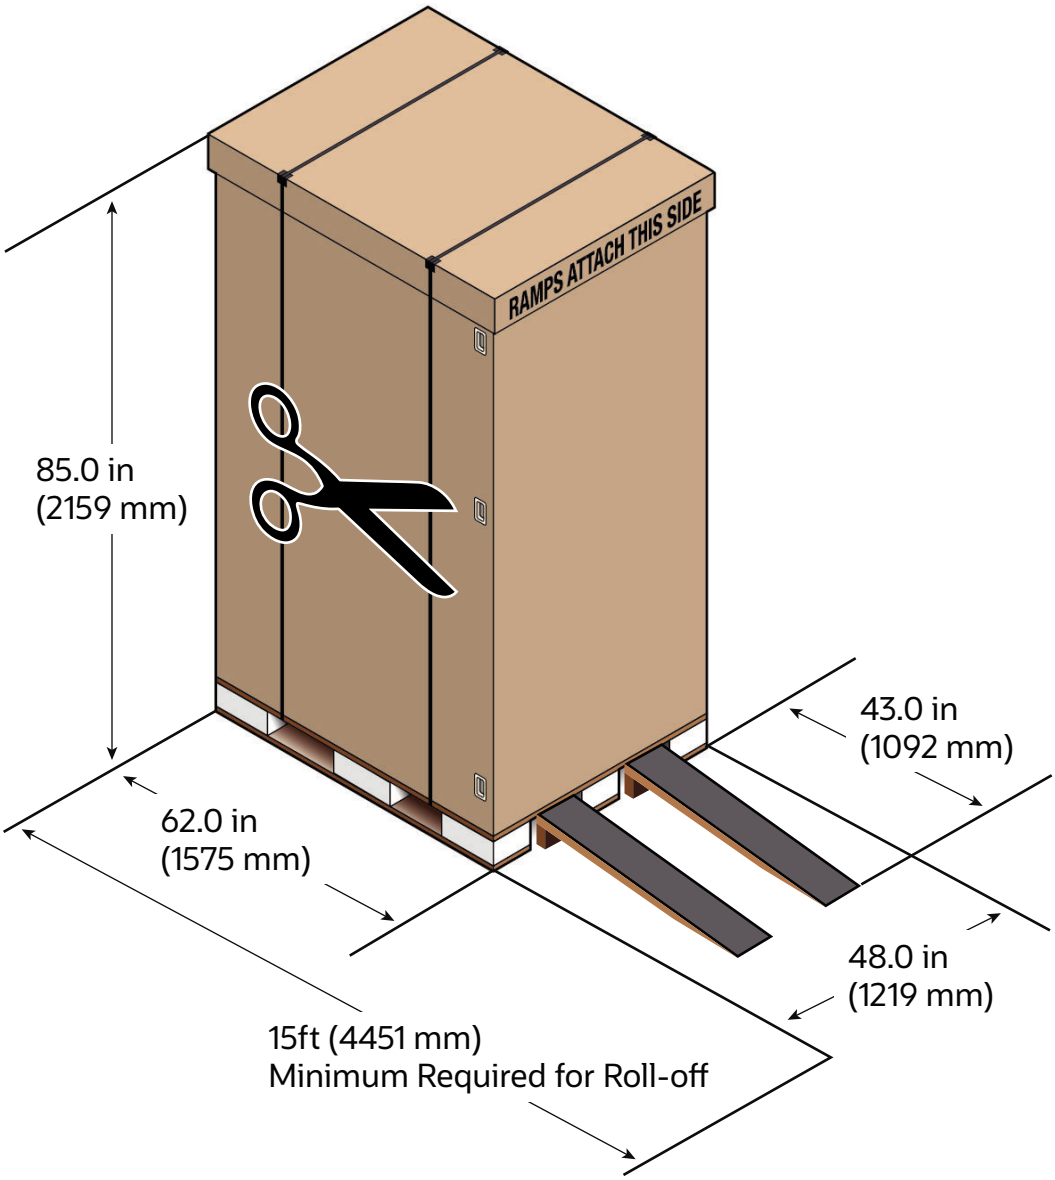

Contents of the Shipping Kit for X6 and Earlier Racks

The ZDLRA Rack shipping kit contains the following tools and equipment to install and service the Engineered System rack:

- 16 mm long No. 2 Phillips screw
- T30 Torx cranked wrench key
- T25 Torx cranked wrench key
- 6 mm hexagon Allen wrench key
- SW 12 mm single-headed wrench
- 2 square jumper brackets with 4 M5 Torx screws

- 2 cable management hooks with 4 spring nuts
- Side panel removal tool
- Keys to the front door, rear door, and side panel locks
- 32 M6 cage nuts
- 32 M6 screws
- Cage nut mounting tool
- SW 17 mm single-headed wrench is included on the shipping pallet

> **Note:**
>
> The following items are not included in the shipping kit:
>
> - No. 2 Phillips screwdriver
> - Antistatic wrist strap

Removing ZDLRA Rack from the Shipping Crate

Use these steps to unpack ZDLRA Rack and prepare it for moving to the installation site.

> **WARNING:**
>
> Oracle strongly recommends that you use professional movers when unpacking and installing ZDLRA Rack. Rocking or tilting the rack can cause it to fall over and cause serious personal injury or death.

1. Unpack ZDLRA Rack carefully from the packaging and shipping pallet as follows:

   > **Note:**
   >
   > Before unpacking the rack from the shipping carton, refer to the labels on the carton and to the instructions that they provide.

   a. Remove the shipping carton bands.
   b. Remove the carton top.
   c. Remove the carton sides and inner top.
2. Remove the shipping kit.
3. Attach the ramps to the shipping pallet as follows:
   a. Remove the ramps from the pallet sides.
   b. Obtain the parts bag from inside the cabinet.

c. Adjust the leveling bolts on the ramps and connect the ramps to the pallet wheel track.

d. Verify that leveling feet are up before moving rack.

4. Carefully roll ZDLRA Rack off the shipping pallet as follows:

   a. Unfasten the exterior shipping brackets from the pallet.

   b. Unfasten the interior shipping brackets from the pallet.

> **Note:**
>
> Use care when removing the shipping brackets from underneath ZDLRA Rack. Access to the inside shipping brackets might be limited.

   c. Roll the cabinet down the ramps to the level floor. Oracle recommends having three people available to move the rack down the ramp: two people on both sides to help guide the rack and one person in back.

      The rack's front casters are fixed, so you must steer using the rear casters. You can maneuver the rack safely by pushing it from behind.

5. Save the shipping brackets used to secure the rack to the shipping pallet. Do not dispose of these brackets, because you cannot order replacement brackets.

6. Recycle the packaging properly. Follow local laws and guidelines to dispose of the material.

> **See Also:**
>
> - *Sun Rack II User's Guide* at `https://docs.oracle.com/cd/E19657-01/html/E29153/index.html`
> - *Oracle Rack Cabinet 1242 User's Guide* at `http://docs.oracle.com/cd/E85660_01/html/E87280/index.html` for X7 or later systems

Placing ZDLRA Rack in Its Allocated Space

This section describes how to position, stabilize, and ground ZDLRA Rack.

Moving ZDLRA Rack

The following procedure describes how to move ZDLRA Rack:

1. Ensure the doors are closed and secured.
2. Ensure the leveling and stabilizing feet on the rack are raised and out of the way.
3. Push ZDLRA Rack from behind to the installation site.

   When moving ZDLRA Rack to the installation site, the front casters do not turn; you must steer the unit by moving the rear casters. You can safely maneuver ZDLRA Rack by carefully pushing it. Figure 5-2 shows the correct way to push the rack.

Figure 5-2 Carefully Push ZDLRA Rack from Behind

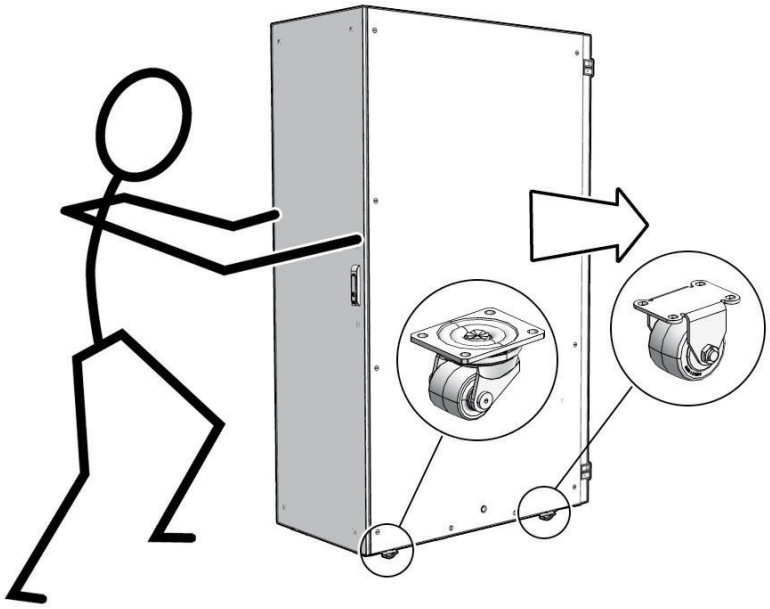

Oracle recommends having two people to move the rack: one person in front and one person in back to help guide the rack. When transporting configured racks from one location to another, take care to move them slowly, 0.65 meters per second (2.13 feet per second) or slower.

Carefully examine the transportation path. Avoid obstacles such as doorways or elevator thresholds that can cause abrupt stops or shocks. Go around obstacles by using ramps or lifts to enable smooth transport.

WARNING:

- Never attempt to move ZDLRA Rack by pushing on the side panels. Pushing on the side panels can tip the rack over. This can cause serious personal injury or death as well as damage to the equipment.
- Never tip or rock ZDLRA Rack because the rack can fall over.

Stabilizing ZDLRA Rack

After moving ZDLRA Rack to the installation site, stabilize the rack to ensure that it does not move or tip over.

Oracle does not recommend using the shipping brackets to permanently mount the rack to the floor.

> **Caution:**
>
> Shipping brackets are not for use for bracing or anchoring the rack during seismic events.

You can stabilize the rack by extending the rack leveling feet.

- Stabilize ZDLRA Rack with Leveling Feet
- Stabilize ZDLRA Rack X7 and Later with Leveling Feet

Stabilize ZDLRA Rack with Leveling Feet

The ZDLRA Rack contains four leveling feet that can be lowered to stabilize the rack. The leveling feet can be used even when the rack is permanently secured to the floor. The four leveling feet should share the rack load with at least one other support. The following is an acceptable method to share the rack load:

- Casters and leveling feet

To adjust the leveling feet, do the following:

1. Locate the four leveling feet located at the bottom corners of the ZDLRA Rack. Figure 5-3 shows the location of the leveling feet on the bottom of the rack.

Figure 5-3 Location of Leveling Feet on ZDLRA Rack

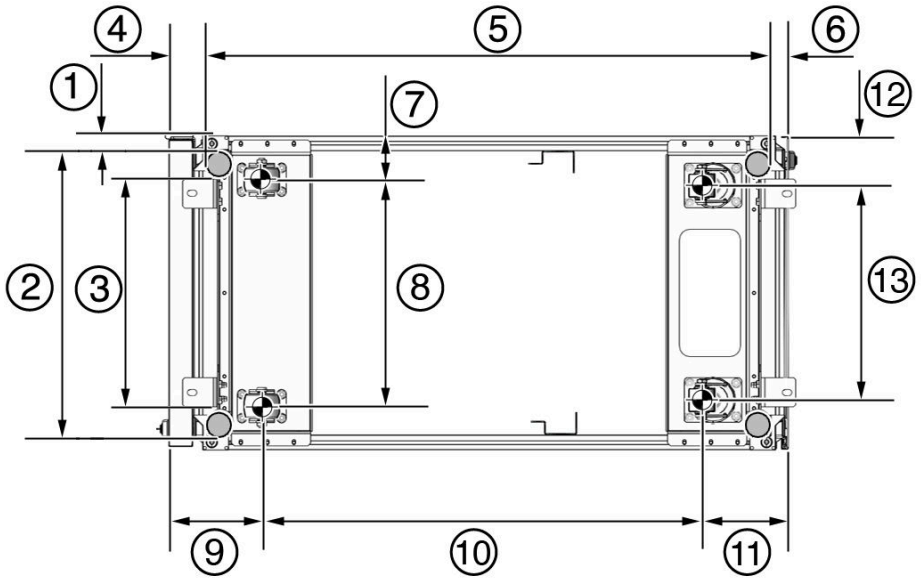

- 1: Distance from the edge of the mounting feet to the side of the rack is 33.75 mm (1.33 inches)
- 2: Width from the outside edges of the leveling feet is 532.5 mm (20.96 inches)
- 3: Width from the inside edges of the leveling feet is 429 mm (16.89 inches)
- 4: Distance from the edge of the feet to the front rack surface is 73.75 m (2.90 inches)

- 5: Depth of the outside edges of the leveling feet is 1058.5 mm (41.67 inches)
- 6: Distance from the edge of the leveling feet to the rear rack surface is 32.5 mm (1.28 inches)
- 7: Distance from the center of front casters to the side of the rack is 86.7 mm (3.41 inches)
- 8: Width between the center of the front casters is 426.6 mm (16.80 inches)
- 9: Distance from the center of the rear casters to the rear of the rack is 173.7 mm (6.83 inches)
- 10: Depth between the front and rear casters is 828.6 mm (32.62 inches)
- 11: Distance between the rear casters and the rear of the rack is 162.4 mm (6.39 inches)
- 12: Distance from the center of rear casters to the side of the rack is 96.4 mm (3.80 inches)
- 13: Width between the center of the rear casters is 407.2 mm (16.03 inches)

2. Lower the leveling feet to the floor to share the load with the casters, as shown in Figure 5-4 using the SW 12 mm wrench.

Figure 5-4 Securing ZDLRA Rack Using the Leveling Feet

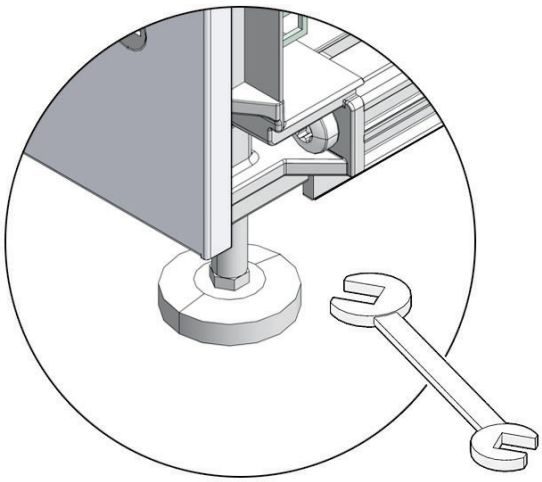

Stabilize ZDLRA Rack X7 and Later with Leveling Feet

The rack contains four leveling feet that can be lowered to stabilize the rack.

The leveling feet can be used even when the rack is permanently secured to the floor. The four leveling feet share the load with the casters. This increases the footprint of the rack which improves stability and helps prevent rack movement.

To adjust the leveling feet, do the following:

1. Locate the four leveling feet located at the bottom corners of ZDLRA Rack. The following figure shows the location of the leveling feet on the bottom of ZDLRA Rack. The image dimensions include the doors attached to the rack.

Figure 5-5 Location of Leveling Feet on ZDLRA Rack X7 or later

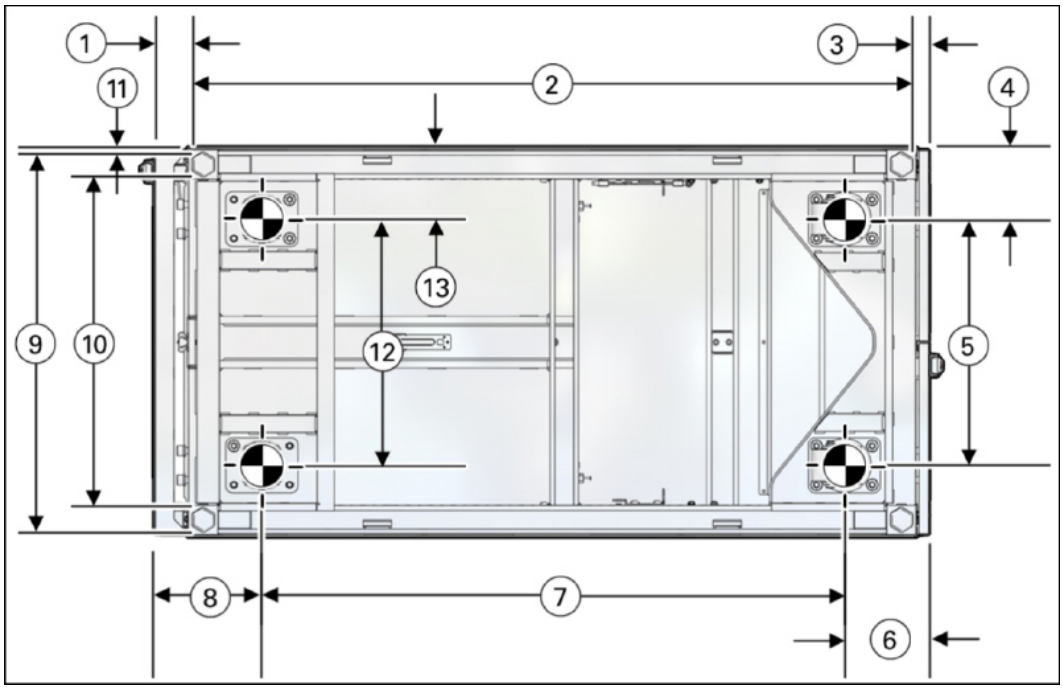

- 1: Distance from the edge of the feet to the Front door surface is 58 mm (2.3 inches)
- 2: Depth of the outside edges of the leveling feet is 1113 mm (43.8 inches)
- 3: Distance from the edge of the leveling feet to the rear door surface is 26 mm (1.0 inches)
- 4: Distance from the center of rear casters to the side of the rack is 96.5 mm (3.8 inches)
- 5: Width between the center of the rear casters is 407 mm (16.0 inches)
- 6: Distance from the center of the rear casters to the rear door surface is 142 mm (5.6 inches)
- 7: Depth between the front and rear casters is 889 mm (35.0 inches)
- 8: Distance between the center of the front casters and the Front door surface is 166 mm (6.5 inches)
- 9: Width from the outside edges of the leveling feet is 583 mm (23.0 inches)
- 10: Width from the inside edges of the leveling feet is 513 mm (20.2 inches)
- 11: Distance from the edge of the mounting feet to the side of the rack is 8.5 mm (0.3 inches)
- 12: Width between the center of the front casters is 410 mm (16.1 inches)
- 13: Distance from the center of front casters to the side of the rack is 95 mm (3.7 inches)

2. Stabilize the rack by lowering the four leveling feet to the floor, using a 6-mm hex wrench or your fingers, as shown in the following image.

Figure 5-6 Securing ZDLRA Rack Using the Leveling Feet

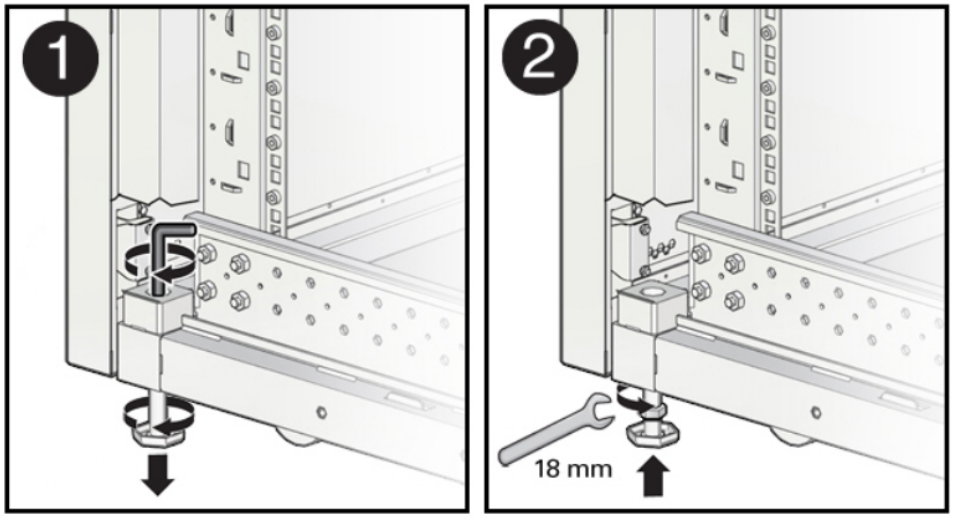

> **See Also:**
>
> *Oracle Rack Cabinet 1242 User's Guide* at `http://docs.oracle.com/cd/E85660_01/html/E87280/index.html` for the rack specifications

Attaching a Ground Cable (Optional)

The ZDLRA Rack power distribution units (PDUs) achieve earth ground through their power cords. Final chassis ground is achieved by way of the ground prong when you connect the power cord to a socket. For additional grounding, attach a chassis earth ground cable to ZDLRA Rack. The additional ground point enables electrical current leakage to dissipate more efficiently.

> **WARNING:**
>
> The PDU power input lead cords and the ground cable must reference a common earth ground. If they do not, then a difference in ground potential can be introduced. If you are unsure of your facility's PDU receptacle grounding, then do not install a ground cable until you confirm that there is a proper PDU receptacle grounding. If a difference in ground potential is apparent, then you must take corrective action.

> **Note:**
>
> A grounding cable is not shipped with the system.

1. Ensure the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.
2. Ensure all grounding points, such as raised floors and power receptacles, reference the facility ground.
3. Ensure that direct, metal-to-metal contact is made for this installation. During manufacturing, the ground cable attachment area might have been painted or coated.
4. Attach the ground cable to one of the attachment points located at the bottom rear of the system frame as shown in Figure 5-7. The attachment point is an adjustable bolt that is inside the rear of the ZDLRA Rack cabinet on the right side.

Figure 5-7 Earth Ground Attachment Bolt Location

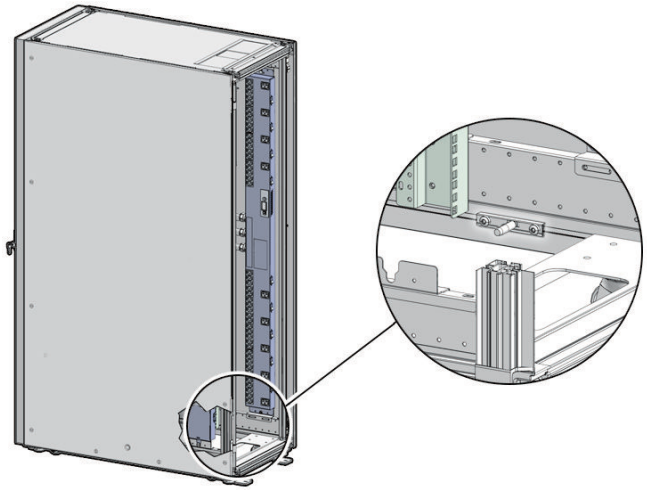

See Also:

Oracle Exadata Database Machine Extending and Multi-Rack Cabling Guide if ZDLRA Rack will be connected to another ZDLRA Rack.

Installing the Fiber Channel Adapters

The optional QLogic fiber channel adapters are shipped separately from the Recovery Appliance rack. Install them only when Recovery Appliance uses an Oracle tape system for backup.

You install a fiber channel adapter in the first PCIe slot of each compute server.

See Also:

Oracle Server X7-2 Service Manual for instructions for installing a PCI card in slot 1:

https://docs.oracle.com/cd/E72435_01/html/E72445/grdej.html#scrolltoc

Oracle Server X6-2 Service Manual for instructions for installing a PCI card in slot 1:

https://docs.oracle.com/cd/E62159_01/html/E62171/z40007611438967.html

Oracle Server X5-2 Service Manual for instructions for installing a PCI card in slot 1:

http://docs.oracle.com/cd/E41059_01/html/E48312/napsm.z40007611438967.html#scrolltoc

Acclimating the Recovery Appliance

Oracle recommends an acclimation period of 24 hours for the Recovery Appliance.

This applies to all hardware, including racks and cables. Acclimation to the surrounding environment prevents damage as a result of condensation.

Powering on the System for the First Time

Before powering on the system for the first time, it is necessary to inspect the rack and connect the power cords. This section contains the following topics:

- Inspecting Recovery Appliance After It Is in Place
- Connecting Power Cords
- Powering Up Recovery Appliance

Note:

After powering on Recovery Appliance, configure it as described in Configuring a Recovery Appliance Rack.

Inspecting Recovery Appliance After It Is in Place

You should examine the Recovery Appliance physical system after it is in place, but before power is supplied.

To inspect the Recovery Appliance rack:

1. Check the rack for damage.

2. Check the rack for loose or missing screws.
3. Check the rack for the ordered configuration. Refer to the Customer Information Sheet (CIS) on the side of the packaging.
4. Check that all cable connections are secure and firmly in place:
   a. Check the power cables. Ensure that the connectors are the correct type for the data center facility power source.
   b. Check the network data cables.
5. Check the site location tile arrangement for cable access and airflow.
6. Check the data center airflow that leads into the front of the rack.

> **See Also:**
>
> "#unique_126"

Connecting Power Cords

> **WARNING:**
>
> Confirm that all breaker switches are off before connecting the power cords. When off (0), the breakers extend beyond the side of the PDU.
>
> There is one switch for each group and six switches for each PDU.

To connect power cords to the Recovery Appliance rack:

1. Open the rear cabinet door.
2. Ensure that the correct power connectors have been supplied.
3. Unfasten the power cord cable ties. The ties are for shipping only and are no longer needed.
4. Route the power cords to the facility receptacles through the bottom, as shown in Figure 5-8.
5. Secure the power cords in bundles as shown in Figure 5-8.
6. Ensure that the breaker switches are OFF before connecting the power cables.
7. Plug the PDU power cord connectors into the facility receptacles.

Figure 5-8 Power Cord Routing Example

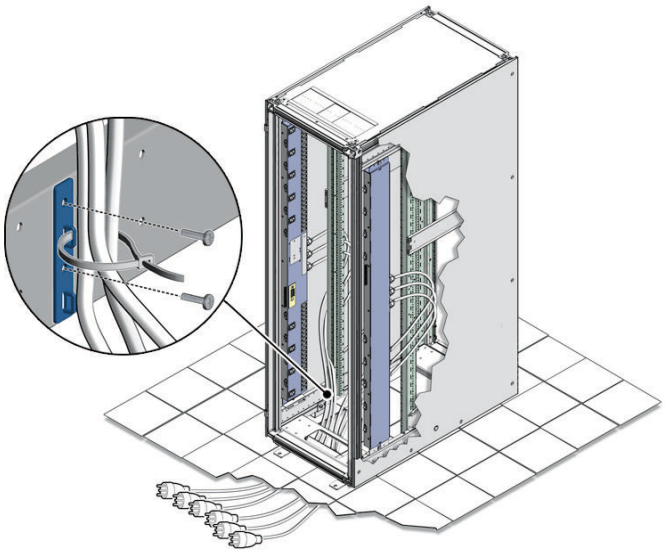

Powering Up Recovery Appliance

The compute servers may take up to 15 minutes to start through the normal BIOS POST tests.

To power up Recovery Appliance:

1. Switch on the PDU B circuit breakers, one at a time. PDU B is on the right side of the rack when viewed from the rear. The circuit breakers are on the rear of the cabinet as shown in Figure 5-9. Press the ON (|) side of the toggle switches, so that the breakers are flush with the side of the PDU.

Figure 5-9 PDU Switch Locations

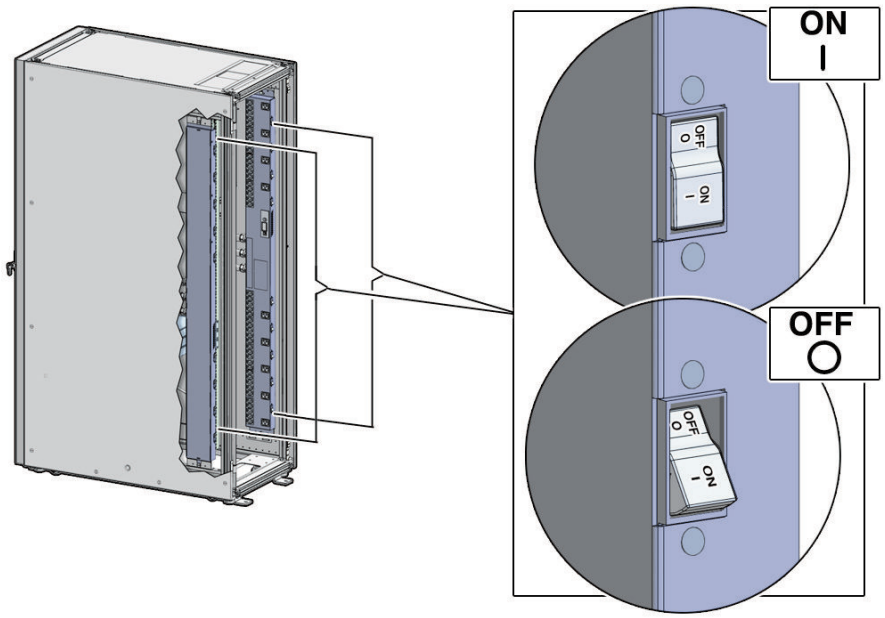

2. Verify that the expected power LEDs are on. The LEDs are located as follows:
   - **Compute servers**: Right LED
   - **Storage servers:** Top LED
   - **Cisco switch (viewed from front)**: Left LED is green, and the other LEDs are red
   - **InfiniBand switches (viewed from front)**: Left LED labeled PS0, or right LED labeled PS1
3. Connect power to PDU A. Ensure the breaker switches are in the OFF position before connecting the power cables.

   > **See Also:**
   >
   > "Reviewing Safety Guidelines" about single phase PDUs and cabling

4. Switch on the PDU A circuit breakers, one at a time.
5. Verify that the expected LEDs are on.
6. Perform a visual check of all cable connections in the rack. Do not press every connector to verify a connection.
7. Verify that the OK LED is blinking standby for all systems. The OK LED blinks every 3 seconds when in standby mode. It is dark for 2 to 3 minutes while Oracle ILOM is booting before going into standby mode.

6

Expanding a Recovery Appliance Rack with Additional Storage Servers

The Recovery Appliance racks have a flexible configuration that allows a minimum of three storage servers. You can extend a rack up to a total of 18 storage servers for Recovery Appliance X9M, X8M, and X8 through X5. Recovery ApplianceX4 only supports 14 storage servers. All Recovery Appliance racks have two compute servers, regardless of the number of storage servers.

This chapter describes how to expand a Recovery Appliance rack by adding storage servers. It contains the following sections:

- Preparing to Extend Recovery Appliance
- Precautions and Guidelines
- Adding Storage Servers to a Rack
- Making New Storage Available

See Also:

- In-Rack Cabling Tables for Recovery Appliance X8M
- In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5
- In-Rack Cabling Tables for Recovery Appliance X4

Preparing to Extend Recovery Appliance

Before expanding a rack, review the cabling requirements and collect information about the current rack. This section contains the following topics:

- Estimating RDMA Network Fabric Cable Path Lengths
- Obtaining the Current Configuration Information
- Generating the Network Configuration Files
- Moving the Audit and Diagnostic Files
- Synchronizing Release and Patch Levels
- Checking the Recovery Appliance Rack

Estimating RDMA Network Fabric Cable Path Lengths

Ensure that you have the required number and lengths of the RDMA Network Fabric cables for your installation. Cable paths should be as short as possible. After you calculate the

length of a cable path, consider the following ways that the length might increase. Then select the shortest cable to satisfy the length requirement.

Consider the following:

- Bends in the cable path increase the required length of the cable. A cable rarely travels in a straight line from connector to connector. Bends in the cable path are necessary, and each bend increases the total length.
- Bundling increases the required length of the cables. Bundling causes one or more cables to follow a common path. However, the bend radius is different in different parts of the bundle. If the bundle is large and unorganized, and there are many bends, one cable might bend only on the inner radius, while another cable might bend only on the outer radius. In this situation, the difference in the required cable lengths is substantial.
- If you are routing the RDMA Network Fabric cable under the floor, consider the height of the raised floor when calculating the cable path length.

Obtaining the Current Configuration Information

Use the current configuration of the Recovery Appliance rack to plan patching requirements, configure new IP addresses, and so on.

Log into the first compute server in the rack and collect the following information:

- The `exachk` report for the rack. See My Oracle Support Doc ID 1070954.1 for information about the `exachk` utility.
- Image history information. Use the following command:

  ```
  # dcli -g ~/all_group -l root "imagehistory" > imagehistory.txt
  ```

- Current IP addresses of all storage servers and compute servers. Use the following command:

  ```
  # dcli -g ~/all_group -l root "ifconfig" > ifconfig_all.txt
  ```

- Configuration of the storage servers, disks, flash logs, and IORM plans. Use the following commands:

  ```
  # dcli -g ~/cell_group -l root "cellcli -e list cell detail" >
  cell_detail.txt

  # dcli -g ~/cell_group -l root "cellcli -e list physicaldisk detail" >   \
  physicaldisk_detail.txt

  # dcli -g ~/cell_group -l root "cellcli -e list griddisk attributes
  name,offset,size,status,asmmodestatus,asmdeactivationoutcome" > griddisk.txt

  # dcli -g ~/cell_group -l root "cellcli -e list flashcache detail" >
  fc_detail.txt

  # dcli -g ~/cell_group -l root "cellcli -e list flashlog detail" >
  fl_detail.txt

  # dcli -g ~/cell_group -l root "cellcli -e list iormplan detail" >
  iorm_detail.txt
  ```

- HugePages memory configuration on the compute servers. Use the following command:

```
dcli -g ~/dbs_group -l root "cat /proc/meminfo | grep 'HugePages'" > hugepages.txt
```

- InfiniBand switch information. Use the following command:

  ```
  # ibswitches > ibswitches.txt
  ```

- Firmware version of the Sun Datacenter InfiniBand Switch 36 switches. Use the `nm2version` command on each switch.
- The following network files on the first compute server in the rack:
  - `/etc/resolv.conf`
  - `/etc/ntp.conf`
  - `/etc/network`
  - `/etc/sysconfig/network-scripts/ifcfg-*`
- Any users, user identifiers, groups, and group identifiers created for cluster-managed services that must be created on the new servers:
  - `/etc/passwd`
  - `/etc/group`
- The current cluster status. Use the following command:

  ```
  # crsctl stat res -t > crs_stat.txt
  ```

- Patch information from the Grid Infrastructure and Oracle homes. Use the following commands as the Grid Infrastructure home owner and the Oracle home owner:

  ```
  /u01/app/oracle/product/11.2.0/dbhome_1/OPatch/opatch lsinventory -oh    \
  GRID_HOME -detail -all_nodes > opatch_grid.txt

  /u01/app/oracle/product/11.2.0/dbhome_1/OPatch/opatch lsinventory -oh    \
  ORACLE_HOME -detail -all_nodes >> opatch_oracle.txt
  ```

  In the preceding commands, *GRID_HOME* is the path of the Grid Infrastructure home directory, and *ORACLE_HOME* is the path of the Oracle home directory.

Generating the Network Configuration Files

Use the Oracle Exadata Deployment Assistant to list the IP addresses for the new servers. Specify only the number of new servers. Define all other items, such as the Oracle home location and owner, the same as the existing rack configuration. In particular, the bonding configuration of the new servers must match the existing servers in the rack. The Oracle Exadata Deployment Assistant InfiniBand Network Configuration Page enables you to select the type of bonding.

You must install the Recovery Appliance software in the new storage servers, after installing them in the rack.

See Also:

- Using the Deployment Assistant
- Installing the Recovery Appliance Software

Moving the Audit and Diagnostic Files

Move to delete files in the `$GRID_HOME/rdbms/audit` directory and the `$GRID_HOME/log/diagnostics` directory before extending a cluster. Oracle recommends doing this task a day or two before the planned extension to allow sufficient time.

Synchronizing Release and Patch Levels

The servers probably have a later release or patch level than the rack where they will be added. You can either update the rack to the later release or reimage the new servers to match the current rack. In either case, ensure that all servers, and all Sun Datacenter InfiniBand Switch 36 switches, are at the same patch level.

Also consider how to handle the Grid Infrastructure and database home releases and bundle patch updates. If you plan to apply the new patches, then Oracle recommends upgrading the existing servers, so that the new servers inherit the releases as part of the extension procedure. Then fewer servers are patched. Apply the patches to the existing servers in advance, thereby reducing the total amount of work required during the extension.

Precautions and Guidelines

This section lists safety precautions and guidelines for handling the equipment. Familiarize yourself with this information before extending the rack:

- General Safety Precautions
- InfiniBand Cable Precautions
- InfiniBand Cable Bundling Guidelines
- Cable Management Arm Guidelines

General Safety Precautions

Read the following safety guidelines before working on a Recovery Appliance rack:

- "Reviewing Safety Guidelines"
- *Important Safety Information for Sun Hardware Systems* shipped with the rack
- *Zero Data Loss Recovery Appliance Safety and Compliance Guide* for safety notices

InfiniBand Cable Precautions

Review the following InfiniBand cable precautions before working with InfiniBand cables:

- Fiber optic InfiniBand cables with laser transceivers must be type Class 1.
- Do not allow any copper core InfiniBand cable to bend to a radius tighter than 127 mm (5 inches). Tight bends can damage the cables internally.

- Do not allow an optical InfiniBand cable to bend to a radius tighter than 85 mm (3.4 inches). Tight bends can damage the cables internally.
- Do not use zip ties to bundle or support InfiniBand cables. The sharp edges of the ties can damage the cables internally. Use hook-and-loop straps.
- Do not subject an InfiniBand cable to extreme tension. Do not pull or drag it. Pulling on an InfiniBand cable can damage it internally.
- Unroll the full length of an InfiniBand cable.
- Do not twist an InfiniBand cable more than one revolution over its entire length. Twisting an InfiniBand cable can damage it internally.
- Do not route InfiniBand cables where they can be stepped on or rolled over. Crushing can damage the cable internally.

InfiniBand Cable Bundling Guidelines

When bundling InfiniBand cables in groups, use hook-and-loop straps to keep the cables organized. If possible, use color-coordinated straps to help identify the cables and their routing. The InfiniBand splitter and 4X copper conductor cables are thick and heavy for their length. Consider the retention strength of the hook-and-loop straps when supporting cables. Bundle as few cables as reasonably possible. Oracle recommends that no more than eight cables be bundled together.

If the InfiniBand cables break free of their straps and fall, then they might break internally from striking the floor or from sudden changes in tension. Place the hook-and-loop straps as close together as reasonably possible, for example, one strap every foot (0.3 m). If a cable breaks free from a strap, then the cable cannot fall far before it is supported by another strap.

Cable Management Arm Guidelines

Review the following cable management arm (CMA) guidelines before routing the cables:

- Remove all required cables from the packaging, and allow cables to acclimate or reach operating temperature, if possible. This improves the ability to manipulate the cables.
- Label both ends of each cable using a label stock that meets the ANSI/TIA/EIA 606-A standard, if possible.
- Begin the installation procedure in ascending order.
- Only slide out one server at a time. Sliding out more than one server can cause cables to drop or cause problems when sliding the servers back.
- Separate the installation by dressing cables with the least stringent bend radius requirements first. The following bend radius requirements are based on EIA/TIA 568-x standards, and may vary from the manufacturer's requirements:
  - CAT5e UTP: 4 x diameter of the cable or 1 inch/25.4 mm minimum bend radius
  - AC power cables: 4 x diameter of the cable or 1 inch/ 25.4 mm minimum bend radius
  - TwinAx: 5 x diameter of the cable or 1.175 inch/33 mm.
  - Quad Small Form-factor Pluggable (QSFP) InfiniBand cable: 6 x diameter of the cable or 2 inch/55 mm.
  - Fiber core cable: 10 x diameter of the cable or 1.22 inch/31.75 mm for a 0.125 cable.
- Install the cables with the best longevity rate first.

Adding Storage Servers to a Rack

The following sections describe how add storage servers to a Recovery Appliance rack:

- Removing the Doors
- Adding New Servers
- Cabling the Storage Servers
- Closing the Rack

> **Note:**
>
> - Apply any patches to the switches and servers in the rack before extending the hardware.
> - You can extend the hardware while Recovery Appliance is online, and with no downtime. However, take extreme care.

Checking the Recovery Appliance Rack

Examine the Recovery Appliance physical systems for any problems:

1. Check the rack for damage.
2. Check the rack for loose or missing screws.
3. Check Recovery Appliance for the ordered hardware configuration.
4. Check that all cable connections are secure and well seated.
5. Check the power cables.
6. Ensure that the correct connectors were supplied for the data center facility power source.
7. Check the network data cables.
8. Check the site location tile arrangement for cable access and airflow.
9. Check the data center airflow into the front of Recovery Appliance.

Unpacking and Labeling the New Storage Servers

The Recovery Appliance racks are supplied with the InfiniBand and Cat5e harness connections for the maximum number of storage servers.

To prepare the new servers for installation:

1. Unpack all components from the packing carton:
   - 1 storage server
   - 2 power cords and labeling kit
   - 1 cable management arm (CMA)

- 2 rack rails
- Component documentation

2. Lay out the cables for the servers.
3. Unroll the cables and stretch them to remove the bends.
4. Apply the cable labels. Oracle recommends labeling all cables before installation.

> **See Also:**
>
> "Parts for InfiniBand Switches" for information about cable labels

Figure 6-1 shows the components of the server expansion kit.

Figure 6-1 Server Expansion Kit Components

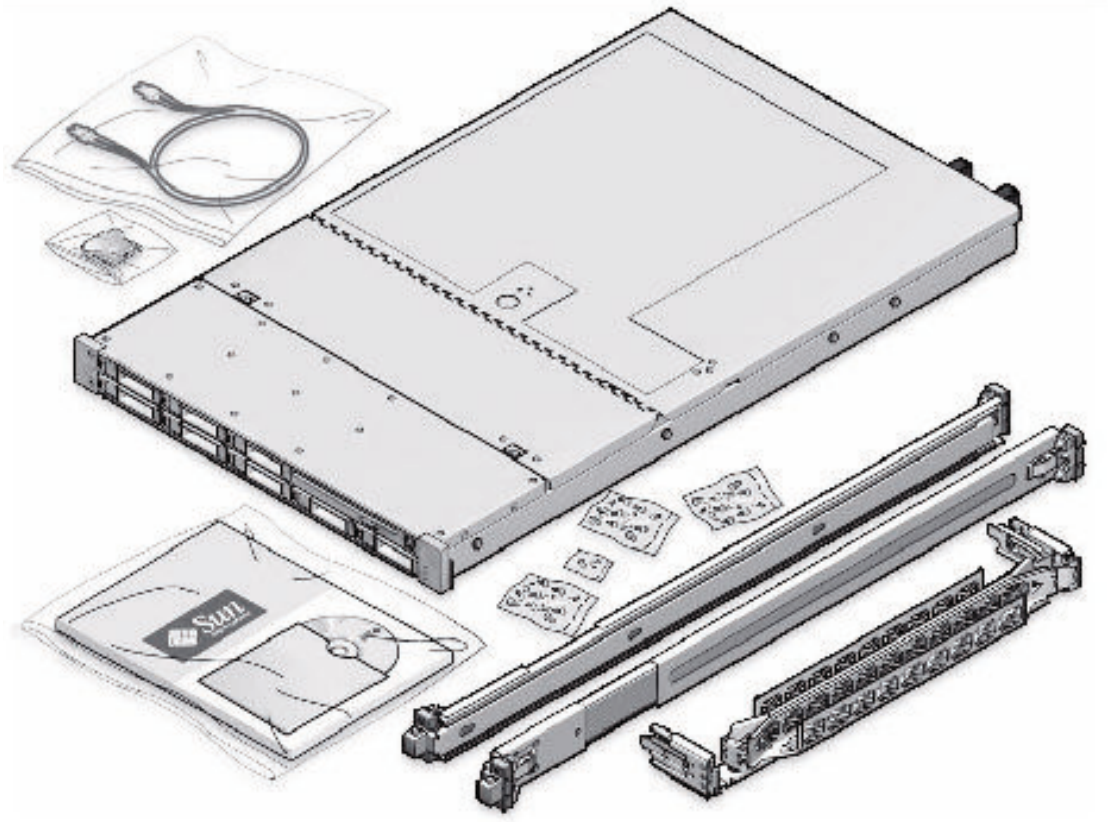

> **Note:**
>
> All new equipment receives a Customer Support Identifier (CSI), including the new equipment for Recovery Appliance. Contact Oracle Support Services to reconcile the new CSI with the existing CSI for the Recovery Appliance rack. You must have the original instance or serial numbers and the new numbers when contacting Oracle Support Services.

Recovery Appliance X9M, X8M, and X8 through X5 Population Rules

When adding servers to a Recovery Appliance X9M, X8M, or X8 through X5 rack, follow the population rules listed in Table 6-1.

Table 6-1 Recovery Appliance X9M, X8M, and X8 through X5 Rack Population Rules

Install Order	Rack Units	Rule
1	U08 to U15	Populate toward the top of the rack (4 units)
2	U27 to U40	Populate toward the top of the rack (7 units)
3	U26 to U23	Populate toward the bottom of the rack (2 units)
4	U19 to U18	(1 unit)
5	U41 to U42	(1 unit)

See Also:

In-Rack Cabling Tables for Recovery Appliance X9M

In-Rack Cabling Tables for Recovery Appliance X8M

In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5

Removing the Doors

To remove the doors on the Recovery Appliance rack:

1. Remove the front and rear doors, as follows:
   a. Unlock the front and rear doors. The keys are in the shipping kit.
   b. Open the doors.
   c. Detach the grounding straps connected to the doors by pressing down on the tabs of its quick-release connectors, and pull the straps from the frame.
   d. Lift the doors up and off their hinges.

Figure 6-2 Removing the Rack Doors

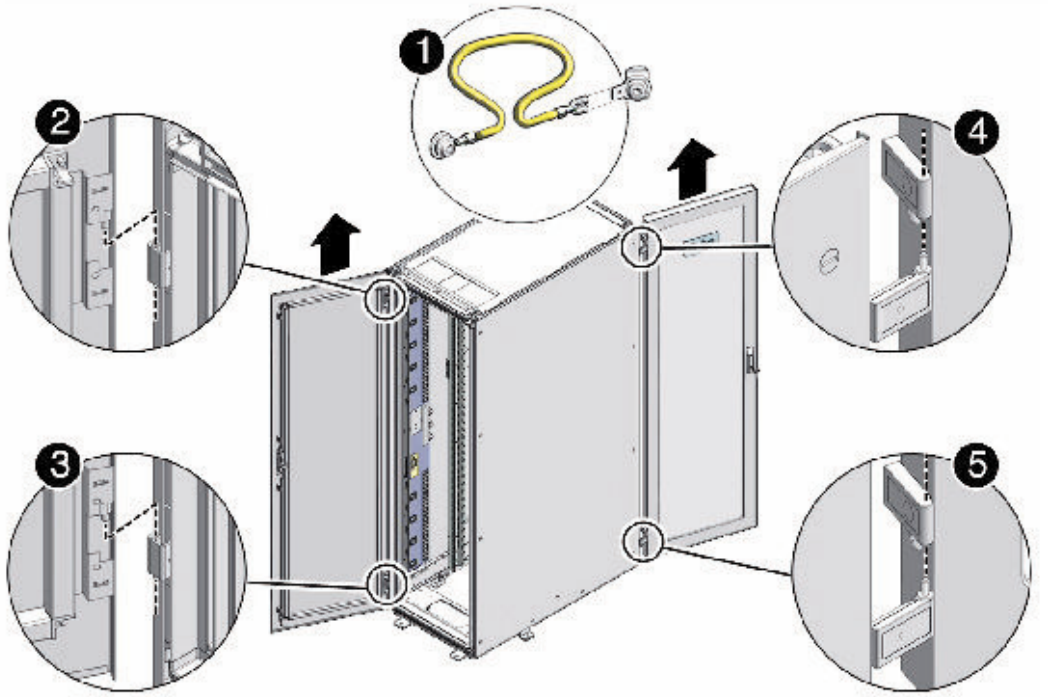

Description of the callouts in Figure 6-2:

1: Detaching the grounding cable

2: Top rear hinge

3: Bottom rear hinge

4: Top front hinge

5: Bottom front hinge

2. To remove the filler panels where the servers will be installed, use a No. 2 screwdriver to remove the M6 screws. The number of screws depends on the type of filler panel. Save the screws for future use.

Figure 6-3 Removing the Filler Panels

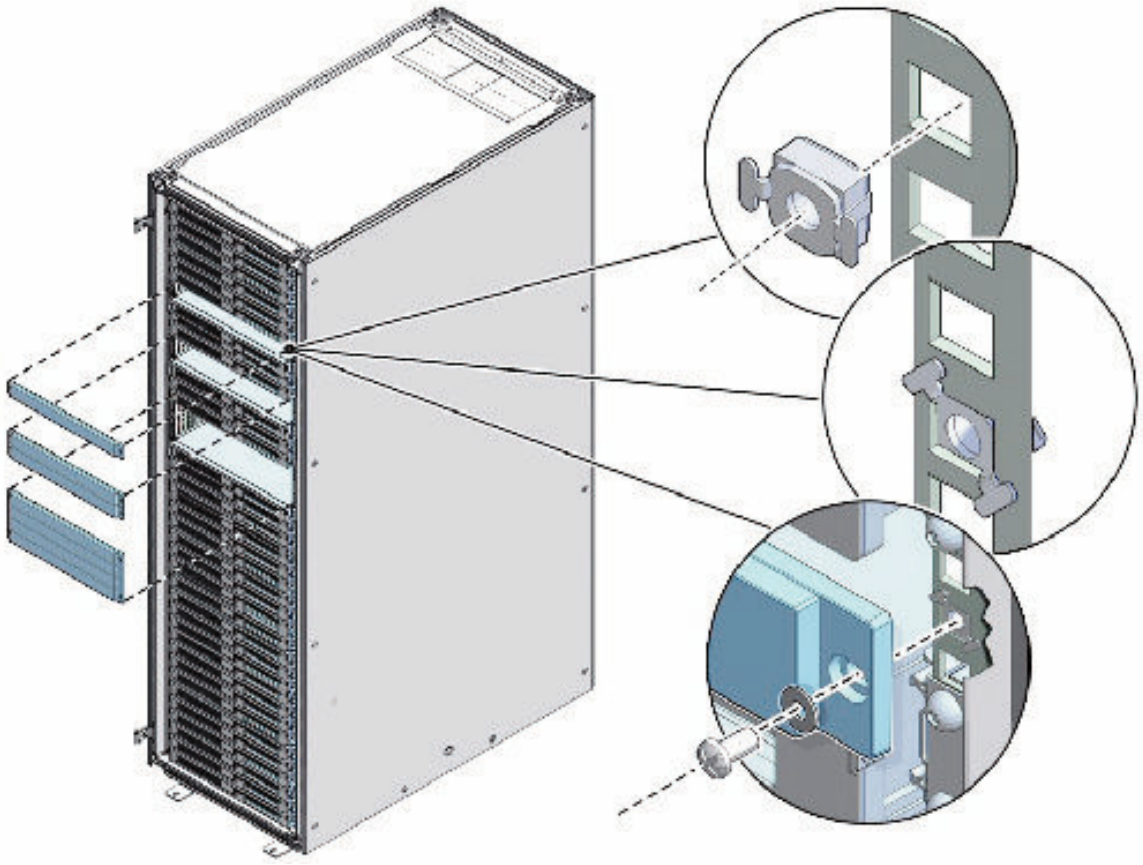

> **Note:**
>
> If you are replacing the filler panels, then do not remove the Duosert cage nuts from the Radio Electronics Television Manufacturers Association (RETMA) rail holes.

Adding New Servers

You can extend Recovery Appliance X7, X6, or X5 to a maximum of 18 storage servers and Recovery Appliance X4 to a maximum of 14 storage servers. You can add any number of storage servers at a time, from one to 15 servers when starting with a minimum configuration.

To prepare for the installation:

1. Identify the rack unit where you will install the server. Fill the first available unit, starting from the bottom of the rack.

   See "Recovery Appliance X8, X7, X6, and X5 Rack Layout" or "Recovery Appliance X4 Rack Layout", as applicable.

2. Remove and discard the trough, which attaches the cable harness when no server is installed in the unit.

Figure 6-4 Location of the Trough

3. Remove and discard the solid filler.

To install the rail assembly:

1. Position a mounting bracket against the chassis so that the slide-rail lock is at the server front, and the five keyhole openings on the mounting bracket are aligned with the five locating pins on the side of the chassis.
2. Orient the slide-rail assembly so that the ball-bearing track is forward and locked in place.
3. Starting on either side of the rack, align the rear of the slide-rail assembly against the inside of the rear rack rail, and push until the assembly locks into place with an audible click.

Figure 6-5 Aligning the Side-Rail Assembly With the Rack

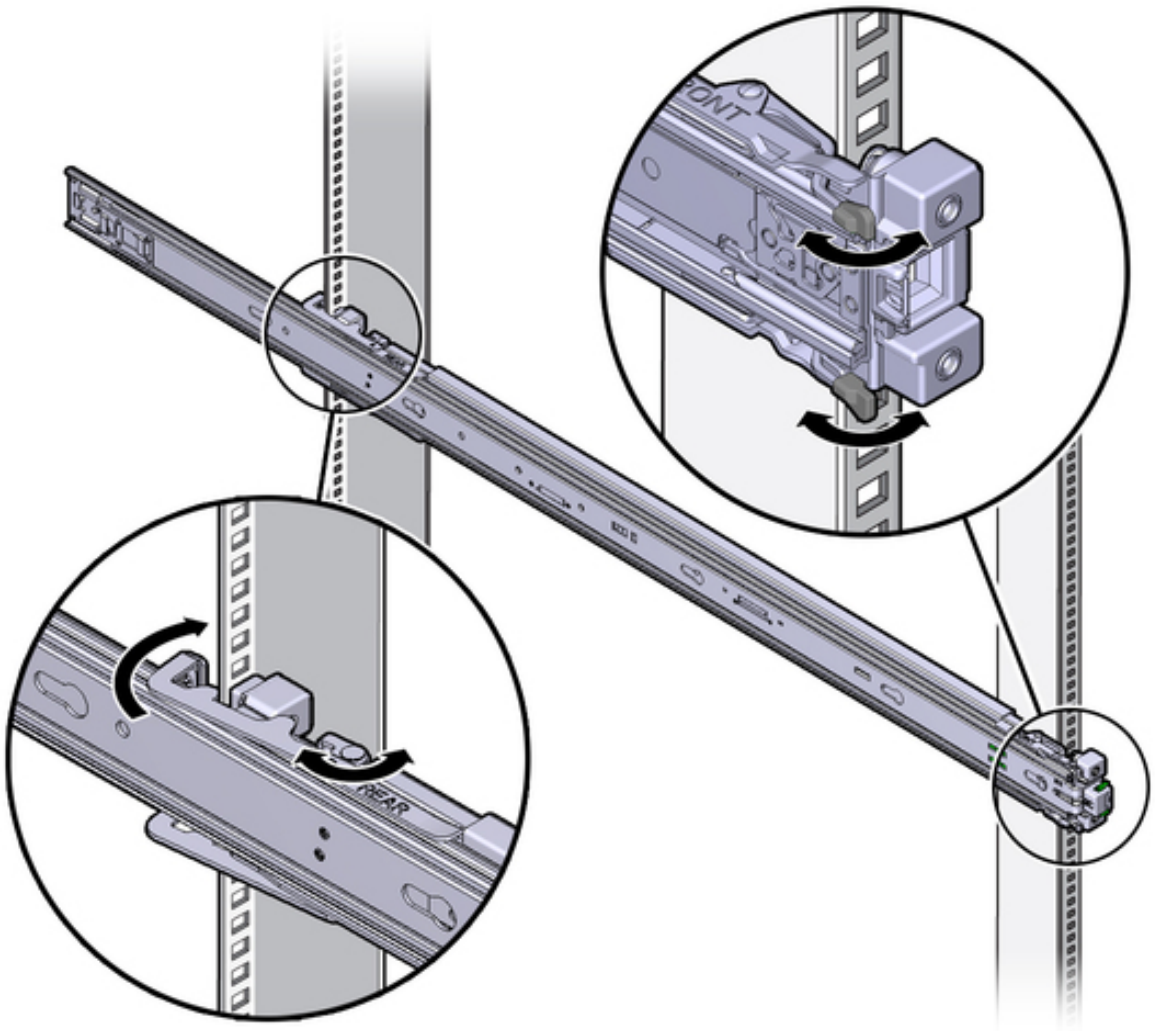

4. Align the front of the slide-rail assembly against the outside of the front rack rail, and push until the assembly locks into place and you hear the click.
5. Repeat steps 2 to 4 on the other side on the rack.

WARNING:

- Installing a server requires a minimum of two people or a lift because of the weight of each server. Attempting this procedure alone can result in equipment damage, personal injury, or both.
- Extend the rack anti-tip bar to prevent the rack from tipping during equipment installation.
- For Recovery Appliance X4 racks, load the servers from the bottom up, so that the rack does not become top-heavy and tip over. See "Recovery Appliance X4 Rack Layout" for the designated location of each server.
- For Recovery Appliance X6 or X5 racks, load the servers following the population rules. See "Recovery Appliance X9M, X8M, and X8 through X5 Population Rules".

To install the server:

1. Read the service label on the top cover of the server before installing a server into the rack.
2. Push the server into the slide rail assembly:
   a. Push the slide rails into the slide rail assemblies as far as possible.
   b. Position the server so the rear ends of the mounting brackets are aligned with the slide rail assemblies mounted in the equipment rack.

Figure 6-6 Inserting the Server with Mounting Brackets into the Slide Rails

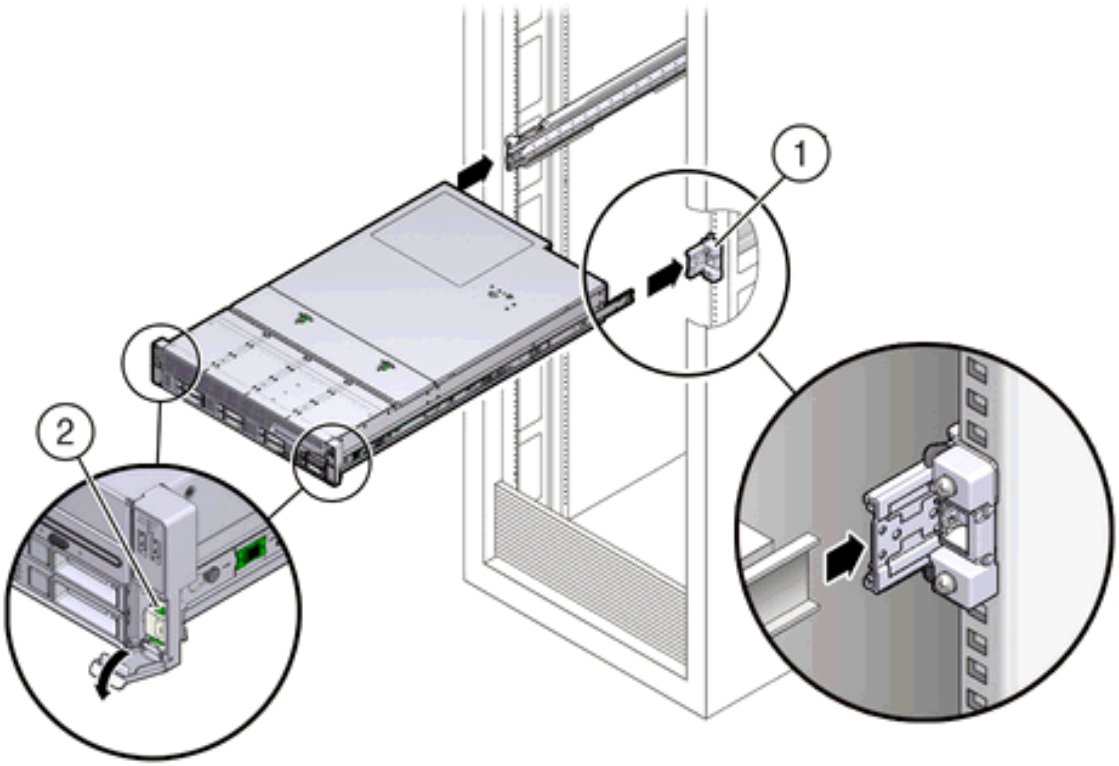

Figure 6-6 callouts highlight the following:

1: Mounting bracket inserted into slide rail

2: Slide-rail release lever

c. Insert the mounting brackets into the slide rails, and push the server into the rack until the mounting brackets encounter the slide rail stops, approximately 30 cm (12 inches).

d. Simultaneously push down and hold the slide rail release levers on each mounting bracket while pushing the server into the rack.

e. Continue pushing until the slide rail locks on the front of the mounting brackets engage the slide rail assemblies, and you hear the click.

3. Cable the new server as described in "Cabling the Storage Servers".

> **Note:**
>
> Oracle recommends that two people push the servers into the rack: one person to move the server in and out of the rack, and another person to watch the movement of the cables and the CMA.

> **See Also:**
>
> "Installing the Server Into a Rack" in *Sun Server X4-2L Installation Guide* at
>
> http://docs.oracle.com/cd/E36974_01/html/E38145/z4000087165586.html#scrolltoc

Cabling the Storage Servers

After installing the new storage servers, cable them to the existing equipment.

> **Note:**
>
> - The existing cable connections in the rack do not change.
> - The blue cables connect to the compute servers, and the black cables connect to the storage servers. These network cables are for the NET0 Ethernet interface port.
> - Attach and route the management cables on the CMA and rear panel for one server at a time. Do not slide out multiple servers.
> - Start from the bottom of the rack, and work upward.
> - Longer hook and loop straps are needed when cabling three CAT5e cables or two twinaxial cables.

To cable the new servers in the rack:

1. Install the cable management arm (CMA):
   a. Into the left slide rail, slide CMA connector A until it locks into place with an audible click.
   b. Into the right slide rail, slide CMA connector B until it locks into place with an audible click. Then do the same with CMA connector C.
2. Insert the cables into their ports through the hook and loop straps, then route the cables into the CMA in this order:
   a. Power
   b. Ethernet
   c. InfiniBand

3. Route the cables through the CMA and secure them with hook and loop straps on both sides of each bend in the CMA.

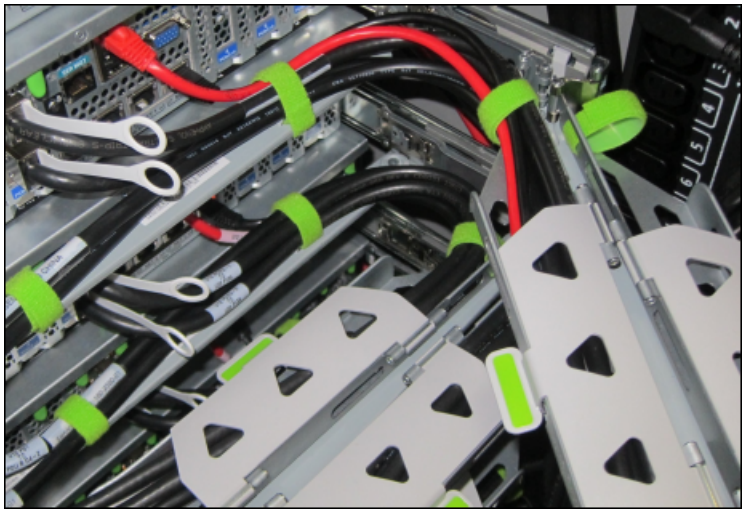

4. Close the crossbar covers to secure the cables in the straightaway.
5. Verify operation of the slide rails and the CMA for each server:

> **Note:**
>
> Oracle recommends that two people do this step: one person to move the server in and out of the rack, and another person to watch the movement of the cables and the CMA.

   a. Slowly pull the server out of the rack until the slide rails reach their stops.
   b. Inspect the attached cables for any binding or kinks.
   c. Verify that the CMA extends fully from the slide rails.
6. Push the server back into the rack:
   a. Release the two sets of slide rail stops.
   b. Locate the levers on the inside of each slide rail, just behind the back panel of the server. They are labeled PUSH.

c. Simultaneously push in both levers and slide the server into the rack, until it stops in approximately 46 cm (18 inches).

d. Verify that the cables and CMA retract without binding.

e. Locate the slide rail release buttons near the front of each mounting bracket.

f. Simultaneously push in both slide rail release buttons and slide the server completely into the rack, until both slide rails engage.

7. Dress the cables, and then tie off the cables with the straps. Oracle recommends that you dress the InfiniBand cables in bundles of eight or fewer.
8. Slide each server out and back fully to ensure that the cables are not binding or catching.
9. Repeat the procedure for all servers.
10. Connect the power cables to the power distribution units (PDUs). Ensure the breaker switches are in the OFF position before connecting the power cables. Do not plug the power cables into the facility receptacles now.

See Also:

- In-Rack Cabling Tables for Recovery Appliance X4
- "Install the Second-Generation Cable Management Arm" in *Sun Server X4-2L Installation Guide* at `http://docs.oracle.com/cd/E36974_01/html/E38145/z40000151427029.html#scrolltoc`

Closing the Rack

To close the rack after installing new equipment:

1. Replace the rack front and rear doors:
   a. Retrieve the doors and place them carefully on the door hinges.
   b. Connect the front and rear door grounding strap to the frame.
   c. Close the doors.
   d. (Optional) Lock the doors. The keys are in the shipping kit.
2. (Optional) Replace the side panels, if they were removed for the upgrade:
   a. Lift each side panel up and onto the side of the rack. The top of the rack can support the weight of the side panel. Ensure that the panel fasteners line up with the grooves in the rack frame.
   b. Turn each side panel fastener 1/4 turn clockwise using the side panel removal tool. Turn the fasteners next to the panel lock clockwise. There are 10 fasteners per side panel.
   c. (Optional) Lock each side panel. The key is in the shipping kit. The locks are located on the bottom center of the side panels.
   d. Connect the grounding straps to the side panels.

After closing the rack:

- If you are installing a new system, proceed to "Installing the Software on Recovery Appliance"
- If you added the server to increase storage on an existing system, proceed to "Making New Storage Available"

Making New Storage Available

After you add storage servers to a rack to increase the storage capacity of an existing system, you must run `racli` commands to make the additional storage available for backups.

> **Caution:**
>
> After you add a storage cell to the Recovery Appliance, you cannot remove it.

On install by default, a DELTA disk group name is created and the storage location name, which is often identical to the disk group name (DELTA, DELTA1, etc.). However the disk group name and storage location name are two different things. Storage is expanded using the UPDATE_STORAGE_LOCATION command.

> **Note:**
>
> The default storage location and the ASM disk group are both called *DELTA*, but are not the same thing.

Before you begin the procedure, ensure that:

- The new storage cells have the same Exadata image version as the existing storage cells.
- The new storage cells do not have any associated databases. (You can run `cellcli -e list database` to verify this is the case.)
- The root password of the new storage cell must still be the factory default.

To make the new storage available:

1. Log in to Recovery Appliance as `root`.
2. Change to the `bin` directory:

   ```
   # cd /opt/oracle.RecoveryAppliance/bin
   ```

3. Run the command to add the new storage cells.

   The following command adds storage cells named `cellhost1` and `cellhost2` to the Recovery Appliance. If three (3) or more cells of different size are being added, then a new disk group DELTA (n) is created, a rebalance is initiated, and expect the rebalance to take more time. The griddisks from the cells are automatically added to the ASM diskgroup.

   ```
   # racli add storage_cells --cells=cellhost1,cellhost2
   ```

4. If you are prompted to do so, run the command to add storage disks.

   The following command adds storage disks with the newly added cells `cellhost1` and `cellhost2`:

   ```
   # racli add storage_disks --cells=cellhost1,cellhost2
   ```

   > **Note:**
   >
   > You can execute this command while the system is running but the operation can take a while to complete and can negatively impact performance. Avoid running other I/O intensive tasks such as patching or maintenance during this time.

5. If you are prompted to update the storage location:

   a. Run the `status` command to ensure that the rebalance operation is complete.

      ```
      # racli status diskgroup
      ```

   b. Run the command to update the storage location.

      The following command increases the size of the storage location using the newly added cells `cellhost1` and `cellhost2`:

      ```
      # racli update storage_location --cells=cellhost1,cellhost2
      ```

6. (Optional) When the procedure is complete, check the size of the storage location.

> **See Also:**
>
> - RACLI Overview for details about the commands used in the procedure.
> - *Zero Data Loss Recovery Appliance Administrator's Guide* for information on how to view storage location information.

7

Extending Recovery Appliance with Multiple Racks

You can extend Recovery Appliance either by adding servers to the current configuration or by cabling together multiple racks.

The procedures and cabling for extending the Recovery Appliance is the same or similar for all engineered systems, and therefore is no longer provided in this manual. This is covered in the *Oracle Engineered Systems Extending and Multi-Rack Cabling Guide* starting with the topic *"*Preparing to Extend the Hardware*"*.

8
Configuring a Recovery Appliance Rack

This chapter describes how to configure the hardware components of a Recovery Appliance rack. It contains the following sections:

> **Note:**
>
> The procedures in this chapter use the files generated by Oracle Exadata Deployment Assistant. You must run this utility before doing the procedures in this chapter.

Supporting Auto Service Request

Auto Service Request is an optional component of Recovery Appliance. To configure Recovery Appliance for Auto Service Request, ASR Manager must be installed first.

Prerequisites for Using Auto Service Request

Verify that Auto Service Request was selected for use in Oracle Exadata Deployment Assistant. Recovery Appliance cannot also be used with Oracle Advanced Support Gateway or Oracle Platinum Gateway.

You must know the IP address and the root password of the ASR Manager host.

Checking an Existing ASR Manager Installation

If ASR Manager is already operating at the site, then verify that it is version 4.5 or higher. Otherwise, you must upgrade it.

To obtain the version number of ASR Manager:

- On a Linux system:

```
# rpm -qa | grep SUNWswasr
SUNWswasr-2.7-1
```

- On a Solaris system:

```
# pkginfo -l SUNWswasr
PKGINST: SUNWswasr
NAME: SASM ASR Plugin
CATEGORY: application
ARCH: all
VERSION: 2.6
BASEDIR: /
VENDOR: Sun Microsystems, Inc.
     .
     .
     .
```

The output from the previous examples indicate that ASR Manager must be updated to 4.5 or higher.

Installing ASR Manager

If ASR Manager is not already installed, then follow the instructions in Setting Up Auto Service Request . After you register ASR Manager with the Oracle ASR back end, return to these instructions for configuring Recovery Appliance.

Installing the Tape Hardware

Oracle Secure Backup tape backup is an option to Recovery Appliance. You must install the QLogic ZLE8362 fiber cards and transceivers on site; they are not factory installed.

The QLogic fiber cards are shipped from Oracle as ride-alongs with the rack. The transceivers are shipped directly from the supplier.

To install the tape networking hardware:

1. Insert a fiber card into PCiE Slot 1 of each compute server.
2. Verify that the cards are seated properly and align with the adjacent cards.
3. Install the transceivers in the SAN switch or tape library, and cable them to the fiber cards.

   If the SAN switch and tape library are not installed yet, then contact your supplier.

> **See Also:**
>
> My Oracle Support Doc ID 1592317.1 for full instructions about replacing a PCIe card

Verifying the Network Configuration Prior to Configuring the Rack

Use the `checkip.sh` script to ensure there are no IP address conflicts between the existing network and your new ZDLRA Rack.

The `checkip.sh` script performs a pre-installation check to verify that the IP addresses and host names that you specified in Oracle Exadata Deployment Assistant (OEDA) are defined in the DNS, that the NTP servers and gateways are available, and that private addresses are not pingable. Running this script before the hardware arrives help to avoid additional delays that would be caused by misconfigured network services, such as Domain Name System (DNS) and NTP.

The `checkip.sh` script is created in a format that matches the operating system of the client on which you ran OEDA. Because this script is run before the engineered system rack has arrived, you typically do not run this script on an engineered system server, but on a client. The client must have access to the same network where the engineered system will be deployed. The script is also available in the ZIP file generated by OEDA.

1. On the client where OEDA was run, copy the `checkip.sh` script generated by OEDA and the XML file *`CustomerName_hostname`*`.xml` to the same directory (one directory level up) as the OEDA `config.sh` script.
2. Run the `checkip.sh` script on the client machine or existing server.

   Use a command similar to the following, where *configuration_file* is the name of the configuration generated by the Oracle Exadata Deployment Assistant for the rack being installed.

   ```
   # ./checkip.sh -cf configuration_file
   ```

   If the command is run from a Microsoft Windows machine, then the command is `checkip.cmd`.
   If this engineered system rack is an addition for an existing installation, then run the `checkip.sh` script from an existing engineered system server. This enables the script to identify in-use IP addresses in the fabric. Not identifying existing IP addresses may cause IP collisions after installation of the new engineered system rack. To create a `checkip.sh` that can run on an existing server, you must run OEDA on a server or client that uses the same operating system as the existing engineered system server.

   The output from the script is a file that contains status messages such as `GOOD` or `ERROR`.

If there are conflicts that you are unable to resolve, then work with your assigned Oracle representative to correct the problems.

Configuring the RDMA Network Fabric Switch

You must perform an initial configuration of the RDMA Network Fabric switch.

Configuring the InfiniBand Switches

The two Sun Datacenter InfiniBand Switch 36 leaf switches are identified in Recovery Appliance as `iba` and `ibb`. Complete these configuration procedures for both switches:

Configuring an InfiniBand Switch

The default identifier for leaf switch 1 in U20 is `iba`, and for leaf switch 2 in U22 is `ibb`.

To configure a Sun Datacenter InfiniBand Switch 36 switch:

1. For a one-rack installation, unplug the InfiniBand cable from Port 8B on the InfiniBand leaf switches. Use hook-and-loop tape to hang it out of the way.

   This cable is preinstalled for a connection to the spine switch in a multirack installation. However, in a one-rack installation, the monitoring software might show it as a down link.

2. Connect a serial cable between your laptop and the InfiniBand switch USB serial adapter. Use these terminal settings, as needed:

   ```
   TERM=vt100; export TERM
   ```

3. Log in to Oracle ILOM:

   ```
   localhost: ilom-admin
   password: welcome1
   ```

4. Ensure that the firmware version is 2.1.5-1 or later:

   ```
   -> version
   ```

   If the switch has a lower version than 2.1.5-1, then contact Oracle Support Services.

5. Set the switch host name, without the domain name. The following example assigns the name `ra1sw` to the first gateway switch (`iba`):

   ```
   -> set /SP hostname=ra1sw-iba
   -> show /SP hostname
   /SP
   Properties:
   hostname = ra1sw-iba
   ```

   See the Installation Template for the name of the switch.

6. Set the DNS server and domain names. In the following syntax, *IP_addresses* can have up to three IP addresses, separated by commas, in the preferred search order.

   ```
   -> set /SP/clients/dns auto_dns=enabled
   -> set /SP/clients/dns nameserver=IP_addresses
   -> set /SP/clients/dns searchpath=domain_name
   ```

7. Verify the settings:

   ```
   -> show /SP/clients/dns
   /SP/clients/dns
   Targets:
   Properties:
   auto_dns = enabled
   nameserver = 10.196.23.245, 138.2.202.15
   retries = 1
   ```

```
searchpath = example.com
timeout = 5
     .
     .
     .
```

8. Configure the switch management network settings. In the following commands, *pending_ip*, *pending_gw*, and *pending_nm* are IP addresses defined by the network administrator:

```
-> cd /SP/network
-> set pendingipaddress=pending_ip
-> set pendingipgateway=pending_gw
-> set pendingipnetmask=pending_nm
-> set pendingipdiscovery=static
-> set commitpending=true
```

9. Verify the settings:

```
-> show
/SP/network
Targets:
test
Properties:
commitpending = (Cannot show property)
dhcp_server_ip = none
ipaddress = 10.196.16.152
ipdiscovery = static
ipgateway = 10.196.23.254
ipnetmask = 255.255.248.0
macaddress = 00:E0:4B:38:77:7E
pendingipaddress = 10.196.16.152
pendingipdiscovery = static
pendingipgateway = 10.196.23.254
pendingipnetmask = 255.255.248.0
state = enabled
     .
     .
     .
```

10. If any of the values are wrong, repeat the `set pendingip`*`parameter`* command, and then the `commitpending=true` command.

Setting the Time on an InfiniBand Switch

To set the time on an InfiniBand switch:

1. Set the time zone, using the value shown in the Installation Template. The following commands display the current setting, change the time zone, and verify the new setting:

```
-> show /SP/clock
-> set /SP/clock timezone=zone identifier
-> show /SP/clock
```

   The Oracle Exadata Deployment Assistant generates the Installation Template. See "#unique_171".

2. Set the SP clock to the current time. Use the time format *MMddHHmmCCyy*, indicating the month, day, hour, minute, century, and year. The following commands display the current setting, change the time, and verify the new setting:

```
-> show /SP/clock
-> set /SP/clock datetime=MMddHHmmCCyy
-> show /SP/clock
```

3. Configure NTP. The following commands configure both the primary (1) and the secondary (2) NTP servers:

```
-> set /SP/clients/ntp/server/1 address=IP_address
-> set /SP/clients/ntp/server/2 address=IP_address
-> set /SP/clock usentpserver=enabled
```

> **Note:**
>
> If the network does not use NTP, then configure the first compute server (U16) as an NTP server before you install the software in Installing the Recovery Appliance Software .

4. Verify the IP address of the primary NTP server:

```
-> show /SP/clients/ntp/server/1
/SP/clients/ntp/server/1
   Targets:

   Properties:
      address = 10.204.74.2

   Commands:
      cd
      set
      show
```

5. Verify the IP address of the secondary NTP server:

```
-> show /SP/clients/ntp/server/2
/SP/clients/ntp/server/2
   Targets:

   Properties:
      address = 10.196.16.1
     .
     .
     .
```

6. Verify the time:

```
-> show /SP/clock
/SP/clock
   Targets:

   Properties:
      datetime = Mon Nov 04 11:53:19 2013
      timezone = EST (US/Eastern)
      usentpserver = enabled
     .
     .
     .
```

Setting the Serial Number on a Spine Switch

In a multirack configuration, set the rack master serial number in the ILOM of the spine switch. Skip this procedure when configuring the leaf switches.

To set the serial number on the spine switch:

1. Set the system identifier to 40 characters or fewer:

   ```
   -> set /SP system_identifier="Oracle ZDLRA X5 serial_number"
   ```

   An invalid property value error indicates too many characters.

2. Verify that the value is set:

   ```
   -> show /SP system_identifier
        /SP
          Properties:
            system_identifier = Oracle ZDLRA X5 AK012345678
   ```

Checking the Health of an InfiniBand Switch

To check the health of an InfiniBand switch:

1. Open the fabric management shell:

   ```
   -> show /SYS/Fabric_Mgmt
   NOTE: show on Fabric_Mgmt will launch a restricted Linux shell.
   User can execute switch diagnosis, SM Configuration and IB
   monitoring commands in the shell. To view the list of commands,
   use "help" at rsh prompt.
   Use exit command at rsh prompt to revert back to
   ILOM shell.
   FabMan@hostname->
   ```

   The prompt changes from -> to FabMan@*hostname*->

2. Check the general health of the switch:

   ```
   FabMan@ra1sw-iba-> showunhealthy
   OK - No unhealthy sensors
   ```

3. Check the general environment.

   ```
   FabMan@ra1sw-iba-> env_test
   NM2 Environment test started:
   Starting Voltage test:
   Voltage ECB OK
   Measured 3.3V Main = 3.28 V
   Measured 3.3V Standby = 3.42 V
   Measured 12V = 12.06 V
        .
        .
        .
   ```

   The report should show that fans 1, 2, and 3 are present, and fans 0 and 4 are not present. All OK and Passed results indicate that the environment is normal.

4. Determine the current InfiniBand subnet manager priority of the switch. Leaf switches must have an smpriority of 5, and spine switches must have a smpriority of 8. The sample output shown here indicates the correct priority for a leaf switch.

```
FabMan@ra1sw-iba-> setsmpriority list
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
```

5. If the priority setting is incorrect, then reset it:

   a. Disable the subnet manager:

   ```
   FabMan@ra1sw-iba->disablesm
   Stopping partitiond daemon.             [ OK ]
   Stopping IB Subnet Manager..            [ OK ]
   ```

   b. Reset the priority. This example sets the priority on a leaf switch:

   ```
   FabMan@ra1sw-iba->setsmpriority 5
   Current SM settings:
   smpriority
   5 controlled_handover TRUE
   subnet_prefix 0xfe80000000000000
   ```

   c. Restart the subnet manager:

   ```
   FabMan@ra1sw-iba->enablesm
   Starting IB Subnet Manager.             [ OK ]
   Starting partitiond daemon.             [ OK ]
   ```

6. Log out of the Fabric Management shell and the Oracle ILOM shell:

   ```
   FabMan@ra1sw-iba-> exit
   -> exit
   ```

7. Log in to Linux as root and restart the switch:

   ```
   localhost: root
   password: welcome1
   [root@localhost ~]# reboot
   ```

8. Disconnect your laptop from the InfiniBand switch.
9. Repeat these procedures for the second InfiniBand leaf switch.

Setting a Spine Switch as the Subnet Manager Master

The InfiniBand switch located in rack unit 1 (U1) is the spine switch. Recovery Appliance has a spine switch only when it is connected to another Recovery Appliance. It is not included as a basic component of the rack.

Perform these steps after the racks are cabled together

The spine switch is the Subnet Manager Master for the InfiniBand subnet. The Subnet Manager Master has priority 8.

To verify the priority setting of the spine switch:

1. Log in to the spine switch as the `root` user.
2. Run the `setsmpriority list` command.

   The command should show that `smpriority` has a value of 8. If `smpriority` has a different value, then do the following:

   a. Use the `disablesm` command to stop the Subnet Manager.

   b. Use the `setsmpriority 8` command to set the priority to 8.

c. Use the `enablesm` command to restart the Subnet Manager.

The other two InfiniBand switches are the leaf switches. The leaf switches are located in rack units 20 and 22 (U20 and U22). They are the Standby Subnet Managers with a priority of 5. You can verify the status using the preceding procedure, substituting a value of 5 in the command shown in step 22.b.

To determine the Subnet Manager Master:

1. Log in as the `root` user on any InfiniBand switch.
2. Display the location of the Subnet Manager Master.

```
# getmaster
20100701 11:46:38 OpenSM Master on Switch : 0x0021283a8516a0a0 ports 36 Sun DCS 36
QDR switch ra01sw-ib1.example.com enhanced port 0 lid 1 lmc 0
```

   The preceding output shows the proper configuration. The Subnet Master Manager is running on spine switch ra01sw-ib1.example.com.

If the spine switch is not the Subnet Manager Master, then reset the Subnet Manager Master:

1. Use the `getmaster` command to identify the current location of the Subnet Manager Master.
2. Log in as the `root` user on the leaf switch that is the Subnet Manager Master.
3. Disable Subnet Manager on the switch. The Subnet Manager Master relocates to another switch.

> **See Also:**
>
> "Disable the Subnet Manager" in *Sun Datacenter InfiniBand Switch 36 User's Guide* at
>
> http://docs.oracle.com/cd/E19197-01/835-0784-05/z4001de61813698.html#z40003f12047367

4. Use the `getmaster` command to identify the current location of the Subnet Manager Master. If the spine switch is not Subnet Manager Master, then repeat steps 2 and 3 until the spine switch is the Subnet Manager Master.
5. Enable Subnet Manager on the leaf switches that were disabled during this procedure.

> **See Also:**
>
> "Enable the Subnet Manager" in *Sun Datacenter InfiniBand Switch 36 User's Guide* at
>
> http://docs.oracle.com/cd/E19197-01/835-0784-05/z4001de61707660.html#z40003f12047359

> **Note:**
>
> If the InfiniBand network consists of four or more racks cabled together, then only the spine switches run Subnet Manager. Disable the Subnet Manager on the leaf switches.

Configuring the Cisco Nexus C9336C-FX2 Switch

The RoCE Network Fabric switch supplied with the engineered system rack is minimally configured during installation.

During initial system configuration, you can reset and configure the switch.

1. Connect from the RoCE Network Fabric switch serial console to a laptop or similar device using the available RJ45 cable.
2. Ensure the terminal session is recorded on the laptop by logging the output.

   The output can be used as a reference that the switch has been configured correctly.
3. Power on the switch.
4. Log in as the `admin` user.

   ```
   User Access Verification
   dbm0sw-rocea0 login: admin
   Password: ********
   ```

   > **Note:**
   >
   > If you do not have the password for the `admin` user, then contact Oracle Support Services.
5. Erase the existing configuration.

   ```
   dbm0sw-rocea0# write erase

   Warning: This command will erase the startup-configuration.

   Do you wish to proceed anyway? (y/n)  [n] y
   ```
6. Restart the system so you can perform the automated setup.

   ```
   dbm0sw-rocea0# reload

   This command will reboot the system. (y/n)?  [n] y

   2017 Aug 31 01:09:00 dbm0sw-rocea0 %$ VDC-1 %$ %PLATFORM-2-
   PFM_SYSTEM_RESET: Manual system restart from Command Line Interface
   ```

```
CISCO SWITCH Ver7.59
Device detected on 0:1:2 after 0 msecs
...
```

7. Switch to normal setup and, when asked if you want to enforce secure password standard, enter `no`, then enter a new password for the `admin` user.

```
Running S93thirdparty-script...

Populating conf files for hybrid sysmgr ...
Starting hybrid sysmgr ...
inserting /isan/lib/modules/klm_cisco_nb.o ... done

Abort Auto Provisioning and continue with normal setup ? (yes/no) [n]: yes

         ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: no

  Enter the password for "admin":
  Confirm the password for "admin":
```

8. When the Basic System Configuration Dialog appears, choose to enter the basic configuration dialog.

```
        ---- Basic System Configuration Dialog VDC: 1 ----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Please register Cisco Nexus9000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus9000 devices must be registered to receive
entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

9. In the basic configuration, you can use the default inputs until asked to enter the switch name.

   In this example, the switch has a name of `test123sw-rocea0`.

```
  Create another login account (yes/no) [n]:
  Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
  Enter the switch name : test123sw-rocea0
```

10. Respond yes when asked to configure Out-of-band management configuration, and specify appropriate network addresses when prompted.

```
Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: yes
     Mgmt0 IPv4 address : 100.104.10.21
     Mgmt0 IPv4 netmask : 255.255.248.0
  Configure the default gateway? (yes/no) [y]:
     IPv4 address of the default gateway : 100.104.10.1
```

11. Respond yes when asked to configure advanced IP options.

```
Configure advanced IP options? (yes/no) [n]: yes
```

12. Respond yes when asked to configure static route (this can be changed later).

```
Configure static route? (yes/no) [n]: yes
```

13. Enter the destination prefix and mask, and other values as prompted.

```
   Destination prefix : 10.100.100.0

   Destination prefix mask : 255.255.255.0

   Next hop IPv4 address : 10.100.100.1
```

14. Configure the DNS IPv4 addresses.

```
Configure the DNS IPv4 address? (yes/no) [n]: yes
   DNS IP address: 10.100.100.2
```

15. Skip configuring the default domain name (this will be configured later).

```
Configure the default domain name? (yes/no) [n]: no
```

16. Accept the default responses until asked to configure SSH and the NTP server.

```
Enable the telnet service? (yes/no) [n]: no
Enable the ssh service? (yes/no) [y]: yes
   Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
   Number of rsa key bits <1024-2048> [1024]: 1024

Configure the NTP server? (yes/no) [n]: yes
     NTP server IPv4 address : 10.100.100.3
```

17. Accept the default responses until asked to specify the CoPP system profile. Enter strict.

```
 Configure default interface layer (L3/L2) [L2]:
 Configure default switchport interface state (shut/noshut)
[noshut]:
 Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: strict
```

18. After reviewing the configuration, save the configuration.

```
The following configuration will be applied:
   no password strength-check
   switchname test123sw-rocea0
   ip route 100.104.8.0 255.255.248.0 100.104.10.1
   vrf context management
   ip route 0.0.0.0/0 100.104.10.1
   exit
    no feature telnet
    ssh key rsa 1024 force
    feature ssh
    ntp server 100.104.10.1
    system default switchport
    no system default switchport shutdown
    copp profile strict
   interface mgmt0
   ip address 100.104.10.21 255.255.248.0
   no shutdown

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]: yes

[########################################] 100%
Copy complete.
```

19. Enable the scp server feature on the switch.

```
test123sw-rocea0# feature scp-server
```

20. Save the running configuration to flash.

```
test123sw-rocea0# copy running-config startup-config
[########################################] 100%
Copy complete.
```

21. Apply the golden configuration on the switch.
    - Starting with Oracle Exadata System Software release 20.1.0, use the procedure described in Applying Golden Configuration Settings on RoCE Network Fabric Switches, in *Oracle Exadata Database Machine Maintenance Guide*.
    - Otherwise, use the following procedure to apply the golden configuration on the switch:

    a. Delete the configuration file on the switch for the target configuration.

> **Note:**
>
> If you do not remove the file you are replacing, then when you attempt to overwrite the file you will get a 'permission denied' error.

Log in to the switch, enter configuration mode, then run a command similar to the following:

```
test123sw-rocea0# delete bootflash:roce_leaf_switch.cfg
Do you want to delete "/roce_leaf_switch.cfg" ? (yes/no/abort)
[y] y
test123sw-rocea0#
```

b. Log in to a server that has SSH access to the switch, and contains the latest RDMA Network Fabric patch ZIP file.

   If you do not have the patch ZIP file available, you can download the latest patch for your Oracle Exadata System Software release. Refer to My Oracle Support Doc ID 888828.1 and search for the term 'RDMA network switch' for information about the latest available patches.

c. Unzip the RDMA Network Fabric patch ZIP file and change directories to the location of the `patchmgr` utility.

d. Locate the golden configuration files in the RDMA Network Fabric patch bundle.

   The files are located within the `roce_switch_templates` directory.

   The golden configuration files are as follows:

   - Single rack leaf: `roce_leaf_switch.cfg`
   - Multi-rack leaf: `roce_leaf_switch_multi.cfg`
   - Multi-rack spine: `roce_spine_switch_multi.cfg`
   - Single rack leaf with Secure Fabric support: `roce_sf_leaf_switch.cfg`
   - Multi-rack leaf with Secure Fabric support: `roce_sf_leaf_switch_multi.cfg`
   - Single rack leaf configured with 23 host ports: `roce_leaf_switch_23hosts.cfg`
   - Multi-rack leaf configured with 23 host ports: `roce_leaf_switch_23hosts_multi.cfg`

e. Copy the golden configuration file to the switch.

   In the following example, `100.104.10.21` represents the IP address of the switch you are configuring.

```
# scp roce_leaf_switch.cfg admin@100.104.10.21:/
User Access Verification
Password:
roce_leaf_switch.cfg 100% 23KB 23.5KB/s 00:00
```

f. Apply the golden configuration file on the switch.

   Use the `run-script` command while connected directly to the switch.

```
test123sw-rocea0# run-script bootflash:roce_leaf_switch.cfg |
grep 'none'
```

> **Note:**
>
> This command may take up to 1-2 minutes on a single-rack switch and up to 3-4 minutes on a multi-rack switch.

g. Verify the switch configuration.

Use the `patchmgr` utility on the server that has SSH access to the switch, and contains the latest RDMA Network Fabric patch bundle.

In the following command, `roceswitch.lst` is a file that contains the switch host name or IP address.

```
# ./patchmgr --roceswitches roceswitch.lst --verify-config
```

22. Backup up the switch configuration.

Follow the steps in Backing Up Settings on the ROCE Switch, in *Oracle Exadata Database Machine Maintenance Guide*.

23. Optional: Set the clock, using the same procedure as in Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch.

Configuring the Cisco Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with Recovery Appliance has IPBASEK9-MZ firmware. The switch is minimally configured during installation. These procedures configure the Cisco Ethernet switch into one large virtual LAN.

Configuring the Cisco Catalyst 4948 Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with ZDLRA Rack is minimally configured during installation.

The minimal configuration disables IP routing, and sets the following:

- Host name
- IP address setup
- Subnet mask
- Default gateway
- Domain name
- Name server
- NTP server
- Time
- Time zone

Before configuring the switch, note the following:

- The Cisco Ethernet switch should not be connected until the running configuration has been verified, and any necessary changes have been made by the network administrator.

- The Cisco Ethernet switch should not be connected to the customer network until the IP addresses on all components have been configured in ZDLRA Rack. This is to prevent any duplicate IP address conflicts which are possible due to the default addresses set in the components when shipped.

Note that the Cisco 4948E-F switch supports multiple uplinks to the customer network by utilizing ports 49 - 52. This is a more complicated switch setup due to the redundant connectivity, and should be performed by the customer's network administrator.

The following procedure describes how to configure the Cisco Ethernet switch. Configuration should be done with the network administrator.

1. Connect a serial cable from the Cisco switch console to a laptop or similar device. An Oracle supplied rollover cable is pre-installed on the Cisco serial console port. Obtain the appropriate adapter and connect it at the end of the rollover cable. An Oracle P/N 530-3100 RJ45-DB9 adapter as used on ILOM ports will also work, connected at the end of the network cable.
2. Ensure the terminal session is recorded on the laptop by logging the output. The output can be used as a reference that the switch has been configured correctly. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.

   ```
   Switch con0 is now available
   Press RETURN to get started.
   ```

3. Change to the enable mode.

   ```
   Switch> enable
   Password: ******
   Switch#
   ```

   > **Note:**
   >
   > If you do not have the password, then contact Oracle Support Services.

4. Check the current version on the switch.

   ```
   Switch# show version
   Cisco IOS Software, Catalyst 4500 L3 Switch Software (cat4500e-
   IPBASEK9-M), Version 15.2(3)E2, RELEASE SOFTWARE (fc1)
   Technical Support: http://www.cisco.com/techsupport
   Copyright (c) 1986-2014 by Cisco Systems, Inc.
   Compiled Tue 11-Mar-14 18:28 by prod_rel_team

   ROM: 12.2(44r)SG12
   zdlra1sw-ip uptime is 1 minute
   System returned to ROM by reload
   System image file is "bootflash:cat4500e-ipbasek9-mz.152-3.E2.bin"
   Hobgoblin Revision 22, Fortooine Revision 1.40
   ...

   Configuration register is 0x2102
   ```

```
Switch#
```

The version of the Cisco 4948E-F switch firmware purchased and shipped by Oracle with Recovery Appliance X6 is IPBASEK9-MZ, which includes `telnet` and `ssh` support. Currently the full release version string is `cat4500e-ipbasek9-mz.152-3.E2.bin`.

5. Configure the network for a single VLAN. The following example assumes you are using IPv4 addressing.

```
Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# interface vlan 1
Switch(config-if)# ip address 10.7.7.34 255.255.255.0
Switch(config-if)# end
Switch# *Sep 15 14:12:06.309:%SYS-5-CONFIG_I:Configured from console by
console
Switch# write memory
Building configuration...
Compressed configuration from 2474 bytes to 1066 bytes [OK ]
```

6. If IP routing is required on the switch, then leave the IP routing setting as the default, and configure the default gateway. Replace 10.7.7.1 with the IP address of the gateway for the installation:

```
Switch#configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)#ip route 0.0.0.0 0.0.0.0 10.7.7.1
Switch(config)#end
*Sep 15 14:13:26.013:%SYS-5-CONFIG_I:Configured from console by console
Switch#write memory
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

7. Set the host name of the switch.

   This example sets the name to ra1sw-ip:

   The system host name is used as the prompt name.

8. Configure up to three DNS servers. Replace the domain name and IP addresses used in this example with the values for the installation:

```
ra1sw-ip#configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
ra1sw-ip(config)#ip domain-name example.com
ra1sw-ip(config)#ip name-server 10.7.7.3
ra1sw-ip(config)#ip name-server 198.51.100.5
ra1sw-ip(config)#ip name-server 10.8.160.1
ra1sw-ip(config)#end
*Sep 15 14:26:37.045:%SYS-5-CONFIG_I:Configured from console by console
ra1sw-ip#write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

If you do not have DNS service available, you must still set the domain-name so that you can configure the SSH keys.

9. (Optional) Set the password.

```
ra1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
ra1sw-ip(config)# enable password password
ra1sw-ip(config)# enable secret password
ra1sw-ip(config)# end
ra1sw-ip# write memory
*Sep 15 14:25:05.893:%SYS-5-CONFIG_I:Configured from console by
console
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

10. Verify telnet access is disabled. Telnet is not secure, and should not be enabled unless there is a compelling reason. To enable telnet, set a password. To disable it, remove the password.

```
ra1sw-ip#configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
ra1sw-ip(config)#line vty 0 15
ra1sw-ip(config)#login
% Login disabled on line 1, until 'password' is set
% Login disabled on line 2, until 'password' is set
 ...
% Login disabled on line 16, until 'password' is set
ra1sw-ip(config)#end
```

If the login command returns output as shown above, then telnet access has been disabled. If instead you get a prompt, then telnet access is not yet disabled so should be disabled now.

```
ra1sw-ip(config-line)#no password
ra1sw-ip(config-line)#end
ra1sw-ip#write memory
Building configuration...
Compressed configuration from 3786 bytes to 1468 bytes [OK ]
```

11. To configure a secure shell (SSH) on the Ethernet switch:

```
ra1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ra1sw-ip(config)# crypto key generate rsa
% You already have RSA keys defined named ra1sw-ip.example.com.
% Do you really want to replace them? [yes/no]: yes
Choose the size of the key modulus in the range of 360 to 2048 for
your General Purpose Keys. Choosing a key modulus greater than 512
may take a few minutes.
How many bits in the modulus [512]: 768

% Generating 768 bit RSA keys, keys will be non-exportable...[OK]
ra1sw-ip(config)# username admin password 0 welcome1
```

```
ra1sw-ip(config)# line vty 0 15
ra1sw-ip(config-line)# transport input ssh
ra1sw-ip(config-line)# exit
ra1sw-ip(config)# aaa new-model

ra1sw-ip(config)# ip ssh time-out 60
ra1sw-ip(config)# ip ssh authentication-retries 3
ra1sw-ip(config)# ip ssh version 2
ra1sw-ip(config)# end
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
ra1sw-ip# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes[OK]
```

12. Set the clock and time zone. The switch keeps internal time in Coordinated Universal Time (UTC) format.

    - To use UTC, use the following command:

      ```
      no clock timezone global configuration
      ```

    - To use a time zone, use the following command:

      ```
      clock timezone zone hours-offset [minutes-offset]
      ```

      In the preceding command, *zone* is the time zone to display when standard time in effect, *hours-offset* is the hours offset from UTC, and *minutes-offset* is the minutes offset from UTC.

    - Daylight savings time (or summer time) is disabled by default. To set summer time hours, use the following command:

      ```
      clock summer-time zone recurring [week day monthhh:mm week day month \
      hh:mm[offset]]
      ```

      In the preceding command, *zone* is the time zone to be displayed when summer time is in effect (EDT, for example), *week* is the week of the month (1 to 5 or last), *day* is the day of the week (Sunday, Monday, ...), *month* is the month (January, February, ...), *hh:mm* is the hours and minutes in 24-hour format, and *offset* is the number of minutes to add during summer time. The default offset is 60 minutes.

    - To manually set the clock to any time use the following command, where the time specified is relative to the configured time zone:

      ```
      clock set hh:mm:ss month day year
      ```

      In the preceding command, *hh:mm:ss* is the time in 24-hour format, *day* is the day by date in the month, *month* is the name of the month, and *year* is the 4-digit year.

    The ordering of commands is important when setting the local time and time zone. For example, to set the local time to US Eastern time:

    ```
    ra1sw-ip# configure terminal
    Enter configuration commands,one per line. End with CNTL/Z.
    ra1sw-ip(config)# clock timezone EST -5
    ```

```
ra1sw-ip(config)# clock summer-time EDT recurring
ra1sw-ip(config)# end
ra1sw-ip# clock set 21:00:00 August 09 2018
ra1sw-ip# write memory
Building configuration...
Compressed configuration from 3784 bytes to 1465 bytes [OK ]
ra1sw-ip# show clock
21:00:06.643 EST Mon Aug 9 2018
```

13. After setting the local time zone, you can configure up to two NTP servers. Replace the IP addresses used in this example with the values for the installation:

```
ra1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
ra1sw-ip(config)# ntp server 10.7.7.32 prefer
ra1sw-ip(config)# ntp server 198.51.100.19
ra1sw-ip(config)# end
*Sep 15 14:51:08.665:%SYS-5-CONFIG_I:Configured from console by
console
ra1sw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes [OK ]
ra1sw-ip# show ntp status
<output will vary per network>
     .
ra1sw-ip# show clock
21:00:23.175 EST Mon Aug 9 2018
```

The NTP server is synchronized to local time when you connect the Cisco switch to the network and it has access to NTP.

Symbols that precede the `show clock` display indicate that the time is the following:

- ***** Not authoritative
- **.** Authoritative, but NTP is not synchronized.
- Authoritative (blank space).

14. Verify the Ethernet configuration using the following command:

```
ra1sw-ip# show running-config
Building configuration...
Current configuration : 3923 bytes
!
version 15.2
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
     .
     .
     .
```

> **Note:**
>
> If any setting is incorrect, then repeat the appropriate step. To erase a setting, enter `no` in front of the same command. For example, to erase the default gateway, use the following commands:
>
> ```
> ra1sw-ip#configure terminal
> Enter configuration commands, one per line. End with CNTL/Z.
> ra1sw-ip(config)# no ip default-gateway 10.7.7.1
> ra1sw-ip(config)# end
> ra1sw-ip#
> *Sep 15 14:13:26.013: %SYS-5-CONFIG_I: Configured from console by
> console
> ra1sw-ip(config)# write memory
> Building configuration...
> Compressed configuration from 2502 bytes to 1085 bytes[OK]
> ```

15. Save the current configuration.

    ```
    ra1sw-ip#copy running-config startup-config
    Destination filename [startup-config]?
    Building configuration...
    Compressed configuration from 2654 bytes to 1189 bytes[OK]
    ```

16. Exit from the session using the following command:

    ```
    ra1sw-ip# exit

    ra1sw-ip con0 is now available

    Press RETURN to get started.
    ```

17. Disconnect the cable from the Cisco console.

    The Cisco switch must not be connected to the management network at this stage. The switch will be connected later after Oracle has configured the systems with the necessary IP addresses and you have worked with the field service engineer to make any additional changes necessary for connecting to the network.

18. To check the Cisco switch, attach a laptop computer to port 48, and ping the IP address of the internal management network to check the configuration.

    Do not connect the switch to the management network.

Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch

The Cisco Nexus 93108-1G or 9348 Ethernet switch supplied with ZDLRA Rack is minimally configured during installation.

Note that the Cisco Nexus 93108-1G or 9348 switch supports multiple uplinks to the customer network by utilizing the QSFP+ ports. This is a more complicated switch setup due to the redundant connectivity, and should be performed by the customer's network administrator.

Whether you are configuring the switch for the first time, or configuring a replacement switch, use the following procedures:

Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch

During the initial configuration, you reset the switch and use the Basic System Configuration Dialog to configure the switch.

Before configuring the switch, note the following:

- The Cisco Ethernet switch should not be connected until the running configuration has been verified, and any necessary changes have been made by the network administrator.
- The Cisco Ethernet switch should not be connected to the customer network until the IP addresses on all components have been configured in ZDLRA Rack. This is to prevent any duplicate IP address conflicts which are possible due to the default addresses set in the components when shipped.

Configuration should be done with the network administrator.

1. Connect from the Cisco switch serial console to a laptop or similar device using the available RJ45 cable.
2. Ensure the terminal session is recorded on the laptop by logging the output.

   The output can be used as a reference that the switch has been configured correctly. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.
3. Power on the switch.
4. Log in as the `admin` user.

   ```
   User Access Verification
   exadatax7-adm0 login: admin
   Password: ********
   ```

   > **Note:**
   >
   > If you do not have the password for the `admin` user, then contact Oracle Support Services.

5. Erase the existing configuration.

   ```
   exadatax7-adm0# write erase

   Warning: This command will erase the startup-configuration.

   Do you wish to proceed anyway? (y/n)  [n] y
   ```

6. Restart the system so you can perform the automated setup.

   ```
   exadatax7-adm0# reload
   ```

```
This command will reboot the system. (y/n)?  [n] y

2017 Aug 31 01:09:00 exadatax7-adm0 %$ VDC-1 %$ %PLATFORM-2-
PFM_SYSTEM_RESET: Manual system restart from Command Line Interface


CISCO SWITCH Ver7.59
Device detected on 0:1:2 after 0 msecs
...
```

7. Switch to normal setup and, when asked if you want to enforce secure password standard, enter `no`, then enter a new password for the `admin` user.

```
Running S93thirdparty-script...

Populating conf files for hybrid sysmgr ...
Starting hybrid sysmgr ...
inserting /isan/lib/modules/klm_cisco_nb.o ... done

Abort Auto Provisioning and continue with normal setup ? (yes/no) [n]: yes

         ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: no

  Enter the password for "admin":
  Confirm the password for "admin":
```

8. When the Basic System Configuration Dialog appears, choose to enter the basic configuration dialog.

```
        ---- Basic System Configuration Dialog VDC: 1 ----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Please register Cisco Nexus9000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus9000 devices must be registered to receive
entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

9. In the basic configuration, you can use the default inputs until asked to enter the switch name.

   In this example, the switch has a name of `test123sw-adm0`.

```
  Create another login account (yes/no) [n]:
  Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
```

```
Enter the switch name : test123sw-adm0
```

10. Respond no when asked to configure Out-of-band management configuration.

```
Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: no
```

11. Respond yes when asked to configure advanced IP options.

```
Configure advanced IP options? (yes/no) [n]: yes
```

12. Respond no when asked to configure static route (this will be configured later).

```
Configure static route? (yes/no) [n]: no
```

13. Enter the destination prefix and mask, and other values as prompted.

```
   Destination prefix : 10.100.100.0

   Destination prefix mask : 255.255.255.0

   Next hop IPv4 address : 10.100.100.1
```

14. Skip configuring the DNS IPv4 addresses (this will be configured later).

```
Configure the DNS IPv4 address? (yes/no) [n]: no
```

15. Skip configuring the default domain name (this will be configured later).

```
Configure the default domain name? (yes/no) [n]: no
```

16. Accept the default responses until asked to configure SSH and the NTP server.

```
Enable the telnet service? (yes/no) [n]: no
Enable the ssh service? (yes/no) [y]: yes
   Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
   Number of rsa key bits <1024-2048> [1024]: 1024

Configure the ntp server? (yes/no) [n]: yes
     NTP server IPv4 address : 10.100.100.3
```

17. Accept the default responses until asked to specify the CoPP system profile. Enter lenient.

```
 Configure default interface layer (L3/L2) [L2]:
 Configure default switchport interface state (shut/noshut)
[noshut]:
 Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: lenient
```

18. After reviewing the configuration, save the configuration.

```
The following configuration will be applied:
   no password strength-check
   switchname test123sw-adm0
  ...

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]: yes

[########################################] 100%
Copy complete.
```

19. Add the VLAN 1 IP address.

```
test123sw-adm0(config)# feature interface-vlan
test123sw-adm0(config)# interface vlan 1
test123sw-adm0(config-if)# ip address 10.100.100.110/24
test123sw-adm0(config-if)# no shutdown
test123sw-adm0(config-if)# exit
```

20. Set the spanning tree port type for ports 1-47.

```
test123sw-adm0(config)# interface E1/1-47
test123sw-adm0(config-if)# spanning-tree port type edge
test123sw-adm0(config-if)# exit
```

21. Set `switchport` on all 48 ports and set port 48 to a network port (instead of a host port).

```
test123sw-adm0(config)# interface E1/1-48
test123sw-adm0(config-if)# switchport
test123sw-adm0(config-if)# exit
test123sw-adm0(config)# interface E1/48
test123sw-adm0(config-if)# spanning-tree port type network
test123sw-adm0(config-if)# ip route 0.0.0.0/0 10.100.100.1
```

22. Configure the DNS information.

```
test123sw-adm0(config)# ip domain-name example.com
test123sw-adm0(config)# ip name-server 10.100.100.2
test123sw-adm0(config)# exit
```

23. Save the current configuration.

```
test123sw-adm0# copy running-config startup-config
[########################################] 100%
Copy complete.
```

24. Optional: Set the clock, as described in the next topic.

Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch

After you have performed the initial configuration, you can adjust the time used by the switch.

1. Log in as the admin user.
2. View the current time.

```
test123sw-adm0(config)# show clock
20:44:52.986 UTC Thu Aug 31 2017
Time source is NTP
```

3. Set the timezone appropriately.

```
test123sw-adm0(config)# clock timezone PST -8 0
```

4. View the modified time.

```
test123sw-adm0(config)# show clock
12:46:22.692 PST Thu Aug 31 2017
Time source is NTP
```

5. Save the configuration.

```
test123sw-adm0# copy running-config startup-config
[########################################] 100%
Copy complete.
```

Disabling Spanning Tree on the Ethernet Switch

Spanning tree is enabled by default on Cisco switches. If you add a switch with spanning tree enabled to the network, then you might cause network problems. As a precaution, you can disable spanning tree from the uplink port VLAN before connecting the switch to the network. Alternatively, you can turn on spanning tree protocol with specific protocol settings either before or after connecting to the network.

To disable spanning tree on the uplink port VLAN:

1. Disable spanning tree on the uplink port VLAN:

```
rasw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
rasw-ip(config)# no spanning-tree vlan 1
rasw-ip(config)# end
rasw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes[OK]
```

2. Verify that spanning tree is disabled:

```
rasw-ip# show spanning-tree vlan 1
Spanning tree instance(s) for vlan 1 does not exist.
```

To re-enable spanning tree protocol with the default protocol settings:

- Use the commands shown in this example:

```
ra1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ra1sw-ip(config)# spanning-tree vlan 1
ra1sw-ip(config)# end
ra1sw-ip# write memory
```

> **See Also:**
>
> *Cisco Switch Configuration Guide* to enable spanning tree protocol with the specific protocol settings required by the data center Ethernet network

Configuring the Power Distribution Units

The power distribution units (PDUs) are configured with static IP addresses to connect to the network for monitoring.

Assigning Network Addresses to the PDUs

To configure the PDU network addresses:

1. Use an RS-232 cable to connect your laptop to the Cisco Ethernet switch SER MGT port.
2. Configure your laptop's terminal emulator to use these settings:
   - 9600 baud
   - 8 bit
   - 1 stop bit
   - No parity bit
   - No flow control
3. Log in to the PDU metering unit as the `admin` user with password `welcome1`.

   Change this password after configuring the network.
4. Enter the network settings for the IP address, subnet mask, and default gateway:

```
pducli -> set net_ipv4_dhcp=Off
set OK
pducli -> set net_ipv4_ipaddr=ip_address
set OK
pducli -> set net_ipv4_subnet=subnet_mask
set OK
pducli -> set net_ipv4_gateway=default_gateway
set OK
```

5. (Optional) Configure the PDU with the DNS server IP addresses:

```
pducli -> set net_ipv4_dns1=domain_name_1
set OK
pducli -> set net_ipv4_dns2=domain_name_2
set OK
```

6. Reset the PDU metering unit:

```
pducli -> reset=yes
set OK
```

7. Remove the RS-232 cable from the SER MGT port.
8. Repeat these steps for the second PDU metering unit.

Configuring the PDU System Time Settings

To configure the PDUs:

1. Connect your laptop to the Ethernet switch.
2. Open a browser and connect to the PDU, using its IP address:

   `https://pdu_ip_address`

   Accept the security note. The Metering Overview page is displayed.
3. Click **Net Configuration** in the upper left, and log in as user `admin` with the password `welcome1`.
4. Select the **System Time** tab.
5. Configure Manual Settings with the current date and time, and then click **Submit**.
6. Configure NTP Server Settings, and then click **Submit**:
   - Select the Enable option.
   - Enter an NTP server IP address, which is listed on the Installation Template.
   - Select **Time Zone** from the menu.
7. Select the **PDU Information** tab.
8. Enter these values, and then click **Submit**:
   - **Name**: PDU host name, provided in the Installation Template; for example, ra5sw-pdua0
   - **Product Identifier** (case sensitive): ZDLRA X5
   - **Rack Serial Number**: Serial number similar to AK12345678
   - **Location** (optional): Site identifier
9. On the Metering Overview page, select **Module Info**.
10. Confirm that the firmware version is 2.01 or higher. If it is not, then upgrade the firmware after you finish this procedure.
11. Click **Logout** to log out of the PDU.
12. Repeat these steps for the second PDU metering unit.
13. Disconnect the PDU metering units from the Cisco Ethernet switch, and connect them to the data center management network.

Upgrading the PDU Firmware

If the PDU firmware is out of date, with a version earlier than 2.01, then download and install the current version.

To upgrade the PDU firmware:

1. Download the current firmware for Enhanced PDUs from My Oracle Support to your laptop.
2. Unzip the downloaded file on your laptop.
3. Open a browser and connect to the PDU, using its IP address:

```
https://pdu_ip_address
```

Accept the security note. The Metering Overview page is displayed.

4. Click **Net Configuration** in the upper left, and log in as user `admin` with the password `welcome1`.
5. Select the **Firmware Update** tab.
6. Click **Browse**, and select `MKAPP_V2.x.DL` from the unzipped, downloaded firmware files on your laptop.
7. Click Submit to update the firmware.

   The PDU reboots automatically when the update is complete.
8. Reconnect your browser to the PDU.
9. Click Module Info on the Metering Overview page, and confirm that the firmware was updated successfully.
10. Click **Logout** to log out of the PDU.
11. Repeat these steps for the second PDU metering unit.

Checking the Health of the Compute Servers

To check the two compute servers in U16 and U17:

1. Power on both compute servers if they are no up already, and wait while they initialize the BIOS and load the Linux operating system.
2. Use a serial cable to connect your laptop to the first compute server's serial MGT port.
3. Configure your laptop's terminal emulator to use these settings:
   - 9600 baud
   - 8 bit
   - 1 stop bit
   - No parity bit
   - No handshake
   - No flow control
4. Log in as the `root` user with the `welcome1` password.
   - On the first compute server (which is connected to your laptop), open the Oracle ILOM console, and then log in:

     ```
     -> start /SP/console
     ```
   - On the second compute server, use SSH to log in. The default factory IP address is 192.168.1.109.
5. Verify that the rack master and host serial numbers are set correctly. The first number must match the rack serial number, and the second number must match the SysSN label on the front panel of the server.

   ```
   # ipmitool sunoem cli "show /System" | grep serial
        serial_number = AK12345678
        component_serial_number = 1234NM567H
   ```
6. Verify that the model and rack serial numbers are set correctly:

```
# ipmitool sunoem cli "show /System" | grep model
     model = ZDLRA X5
# ipmitool sunoem cli "show /System" | grep ident
     system_identifier = Oracle Zero Data Loss Recovery Appliance X5
AK12345678
```

7. Verify that the management network is working:

```
# ethtool eth0 | grep det
Link detected: yes
```

8. Verify that the ILOM management network is working:

```
# ipmitool sunoem cli 'show /SP/network' | grep ipadd
ipaddress = 192.168.1.108
pendingipaddress = 192.168.1.108
```

9. Verify that Oracle ILOM can detect the optional QLogic PCIe cards, if they are installed:

```
# ipmitool sunoem cli "show /System/PCI_Devices/Add-on/Device_1"
Connected. Use ^D to exit.
-> show /System/PCI_Devices/Add-on/Device_1
  /System/PCI_Devices/Add-on/Device_1
  Targets:

  Properties:
    part_number = 7101674
    description = Sun Storage 16 Gb Fibre Channel PCIe Universal FC HBA,
                  Qlogic
    location = PCIE1 (PCIe Slot 1)
    pci_vendor_id = 0x1077
    pci_device_id = 0x2031
    pci_subvendor_id = 0x1077
    pci_subdevice_id = 0x024d

  Commands:
    cd
    show

-> Session closed
Disconnected
```

See "Installing the Tape Hardware" for information about the QLogic PCIe cards.

10. Verify that all memory is present (256 GB):

```
# grep MemTotal /proc/meminfo
MemTotal: 264232892 kB
[
```

The value might vary slightly, depending on the BIOS version. However, if the value is smaller, then use the Oracle ILOM event logs to identify the faulty memory.

11. Verify that the four disks are visible, online, and numbered from slot 0 to slot 3:

```
# cd /opt/MegaRAID/MegaCli/
# ./MegaCli64 -Pdlist -a0 | grep "Slot\|Firmware state"
Slot Number: 0
Firmware state: Online, Spun Up
Slot Number: 1
Firmware state: Online, Spun Up
Slot Number: 2
```

```
Firmware state: Online, Spun Up
Slot Number: 3
Firmware state: Online, Spun Up
```

12. Verify that the hardware logical volume is set up correctly. Look for Virtual Disk 0 as RAID5 with four drives and no hot spares:

```
[root@db01 ~]# cd /opt/MegaRAID/MegaCli
[root@db01 MegaCli]# ./MegaCli64 -LdInfo -lAll -a0
Adapter 0 -- Virtual Drive Information:
Virtual Drive: 0 (Target Id: 0)
Name :DBSYS
RAID Level : Primary-5, Secondary-0, RAID Level Qualifier-3
Size : 1.633 TB
Physical Sector Size: 512
Logical Sector Size : 512
VD has Emulated PD : No
Parity Size : 557.861 GB
State : Optimal
Strip Size : 1.0 MB
Number Of Drives : 4
Span Depth : 1
     .
     .
     .
```

13. Verify that the hardware profile is operating correctly:

```
# /opt/oracle.SupportTools/CheckHWnFWProfile
[SUCCESS] The hardware and firmware matches supported profile for
server=ORACLE_SERVER_X5-2
```

The previous output shows correct operations. However, the following response indicates a problem that you must correct before continuing:

```
[WARNING] The hardware and firmware are not supported. See details below
[InfinibandHCAPCIeSlotWidth]
Requires:
x8
Found:
x4
[WARNING] The hardware and firmware are not supported. See details above
```

Use the `--help` argument to review the available options, such as obtaining more detailed output.

14. When connected to the first compute server only:

    a. Verify the IP address of the first compute server:

```
# ifconfig eth0
eth0 Link encap:Ethernet HWaddr 00:10:E0:3C:EA:B0
     inet addr:172.16.2.44 Bcast:172.16.2.255 Mask:255.255.255.0
     inet6 addr: fe80::210:e0ff:fe3c:eab0/64 Scope:Link
     UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
     RX packets:7470193 errors:0 dropped:0 overruns:0 frame:0
     TX packets:4318201 errors:0 dropped:0 overruns:0 carrier:0
     collisions:0 txqueuelen:1000
     RX bytes:872195171 (831.7 MiB) TX bytes:2444529519 (2.2 GiB)
```

    b. Verify the IP address of the second compute server:

```
# ibhosts
Ca : 0x0010e0000159c61c ports 2 "node4 elasticNode 172.16.2.40,172.16.2.40 ETH0"
```

```
Ca : 0x0010e000015a46f0 ports 2 "node10 elasticNode
172.16.2.46,172.16.2.46 ETH0"
Ca : 0x0010e0000159d96c ports 2 "node1 elasticNode
172.16.2.37,172.16.2.37 ETH0"
Ca : 0x0010e0000159c51c ports 2 "node2 elasticNode
172.16.2.38,172.16.2.38 ETH0"
Ca : 0x0010e000015a5710 ports 2 "node8 elasticNode
172.16.2.44,172.16.2.44 ETH0"
```

15. Disconnect from the server:
    - First compute server: `exit`
    - Second compute server: `logout`
16. Repeat these steps for the second compute server.

Checking the Health of the Storage Servers

A Recovery Appliance X5 and higher versions have three to 18 storage servers, and a Recovery Appliance X4 rack has three to 14 storage servers. Begin at the bottom of the rack and check each server.

To check a storage server:

1. Power on all storage servers if they are not already on, and wait while the servers initialize the BIOS and load the Linux operating system.
2. Use SSH to connect your laptop to the first storage server. Use its factory IP address.

   See "Factory IP Address Settings".
3. Log in as the `root` user with the `welcome1` password.

   The terminal emulation settings are the same as for the compute servers. See "Checking the Health of the Compute Servers".
4. Verify that the rack master and host serial numbers are set correctly. The first number must match the rack serial number, and the second number must match the SysSN label on the front panel of the server.

   ```
   # ipmitool sunoem cli "show /System" | grep serial
        serial_number = AK01234567
        component_serial_number = 1234NM5678
   ```
5. Verify that the model and rack serial numbers are set correctly:

   ```
   # ipmitool sunoem cli "show /System" | grep model
        model = ZDLRA X5
   # ipmitool sunoem cli "show /System" | grep ident
        system_identifier = Oracle Zero Data Loss Recovery Appliance X5
   AK01234567
   ```
6. Verify that the management network is working:

   ```
   # ethtool eth0 | grep det
   Link detected: yes
   ```
7. Verify that the ILOM management network is working:

   ```
   # ipmitool sunoem cli 'show /SP/network' | grep ipadd
   ipaddress = 192.168.1.101
   pendingipaddress = 192.168.1.101
   ```

8. Verify that all memory is present. X5 has 96 GB, while X8 has 384 GB:

```
# grep MemTotal /proc/meminfo
MemTotal: 98757064 kB
[
```

If the value is smaller, then use the Oracle ILOM event logs to identify the faulty memory.

9. Verify that the hardware profile is operating correctly:

```
# /opt/oracle.SupportTools/CheckHWnFWProfile
[SUCCESS] The hardware and firmware matches supported profile for
server=ORACLE_SERVER_X5-2L_EXADATA_HIGHCAPACITY
```

The previous output shows correct operations. However, the following response indicates a problem that you must correct before continuing:

```
[WARNING] The hardware and firmware are not supported. See details below
[InfinibandHCAPCIeSlotWidth]
Requires:
x8
Found:
x4
[WARNING] The hardware and firmware are not supported. See details above
```

Use the `--help` argument to review the available options, such as obtaining more detailed output.

10. Verify that 12 disks are visible, online, and numbered from slot 0 to slot 11:

```
# cd /opt/MegaRAID/MegaCli
# ./MegaCli64 -Pdlist -a0 | grep "Slot\|Firmware state"
Slot Number: 0
Firmware state: Online, Spun Up
Slot Number: 1
Firmware state: Online, Spun Up
     .
     .
     .
```

11. Verify that there are four NVME logical devices:

```
# ls -l /dev | grep nvme | grep brw
brw-rw---- 1 root disk 259, 0 Nov 12 19:10 nvme0n1
brw-rw---- 1 root disk 259, 1 Nov 12 19:10 nvme1n1
brw-rw---- 1 root disk 259, 2 Nov 12 19:10 nvme2n1
brw-rw---- 1 root disk 259, 3 Nov 12 19:10 nvme3n1
```

12. Confirm the healthy status of the AIC card:

```
# nvmecli --identify --all | grep -i indicator
Health Indicator      : Healthy
Health Indicator      : Healthy
Health Indicator      : Healthy
Health Indicator      : Healthy
```

13. Verify that the boot order is USB (Oracle Unigen), RAID, and PXE:

```
[# ubiosconfig export all > /tmp/bios.xml
[# grep -m1 -A20 boot_order /tmp/bios.xml
<boot_order>
  <boot_device>
    <description>USB:USBIN0:ORACLE SSM UNIGEN-UFD PMAP</description>
    <instance>1</instance>
```

```
</boot_device>
<boot_device>
  <description>RAID:PCIE6:(Bus 50 Dev 00)PCI RAID Adapter</description>
  <instance>1</instance>
</boot_device>
<boot_device>
  <description>PXE:NET0:IBA XE Slot 3A00 v2320</description>
  <instance>1</instance>
</boot_device>
<boot_device>
  <description>PXE:NET1:IBA XE Slot 4001 v2196</description>
  <instance>1</instance>
</boot_device>
```

14. If the boot order is wrong, then restart the server and fix the order in the BIOS setup:

    ```
    # ipmitool chassis bootdev bios
    # shutdown -r now
    ```

15. Exit or log out of SSH.
16. Repeat these steps for the next storage server until you have checked all of them.

Verifying the RoCE Network Fabric Configuration

This procedure describes how to verify the RoCE Network Fabric configuration.

1. Verify the proper `oracle-rdma-release` software versions are being used on the database servers.

   ```
   [root@dbm01adm08 ~]# rpm -qa |grep oracle-rdma-release
   oracle-rdma-release-0.11.0-1.el7ora.x86_64
   ```

   The `oracle-rdma-release` software and adapter firmware versions are automatically maintained on the Recovery Appliance storage servers.

2. Check the adapter firmware versions on the database servers.

   Use the `CheckHWnFWProfile` script to check firmware versions for the RDMA Network Fabric adapters.

   ```
   # /opt/oracle.SupportTools/CheckHWnFWProfile -action list
   ```

3. Visually check all the RDMA Network Fabric cable connections within the rack.

   The port lights should be on, and the LEDs should be on. Do not press each connector to verify connectivity.

4. Complete the steps described in My Oracle Support Doc ID 2587717.1

Related Topics

- Verify RoCE Cabling on Oracle Exadata Database Machine X8M-2 and X8M-8 Servers (My Oracle Support Doc ID 2587717.1)

Verifying the InfiniBand Network Fabric Network

This procedure describes how to verify the InfiniBand Network Fabric network.

1. Visually check all the RDMA Network Fabric cable connections within the rack. The port lights should be on, and the LEDs should be on. Do not press each connector to verify connectivity.
2. Log in as the `root` user on any component in the rack.
3. Verify the InfiniBand Network Fabric topology using the following commands:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify-topology [-t rack_size]
```

The following example shows the output when the network components are correct.

```
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch......
[SUCCESS ]
Are any external switches connected to each other................
[SUCCESS ]
Are any hosts connected to spine switch..........................
[SUCCESS ]
Check if all hosts have 2 CAs to different switches..............
[SUCCESS ]
Leaf switch check:cardinality and even distribution..............
[SUCCESS ]
Check if each rack has an valid internal ring....................
[SUCCESS ]
```

In the preceding command, *rack_size* is the size of the rack. The `-t` *rack_size* option is needed if the rack is Recovery Appliance Half Rack or Recovery Appliance Quarter Rack. Use the value `halfrack` or `quarterrack`, if needed.

The following example shows the output when there is a bad RDMA Network Fabric switch to cable connection:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch......
[SUCCESS ]
Are any external switches connected to each other................
[SUCCESS ]
Are any hosts connected to spine switch..........................
[SUCCESS ]
Check if all hosts have 2 CAs to different switches..............
[SUCCESS ]
Leaf switch check:cardinality and even distribution..............
[SUCCESS ]
Check if each rack has an valid internal ring...................[ERROR ]

Switches 0x21283a87cba0a0 0x21283a87b8a0a0 have 6 connections between
them.
They should have at least 7 links between them
```

The following example shows the output when there is a bad RDMA Network Fabric cable on a database server:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch......
[SUCCESS ]
Are any external switches connected to each other................
[SUCCESS ]
Are any hosts connected to spine switch.........................
[SUCCESS ]
Check if all hosts have 2 CAs to different switches..............
[ERROR ]
Node db01 has 1 endpoints.(Should be 2)
Port 2 of this node is not connected to any switch
--------fattree End Point Cabling verification failed-----
Leaf switch check:cardinality and even distribution..............
[ERROR ]
Internal QDR Switch 0x21283a87b8a0a0 has fewer than 4 compute nodes
It has only 3 links belonging to compute nodes
[SUCCESS ]
Check if each rack has an valid internal ring...................
[SUCCESS ]
```

The following example shows the output when there is a bad connection on the switch and the system:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch......
[SUCCESS ]
Are any external switches connected to each other................
[SUCCESS ]
Are any hosts connected to spine switch.........................
[SUCCESS ]
Check if all hosts have 2 CAs to different switches..............
[ERROR ]

Node burxdb01 has 1 endpoints.(Should be 2)
Port 2 of this node is not connected to any switch
--------fattree End Point Cabling verifation failed-----
Leaf switch check:cardinality and even distribution..............
[ERROR ]
Internal QDR Switch 0x21283a87b8a0a0 has fewer than 4 compute nodes
It has only 3 links belonging to compute nodes...................
[SUCCESS ]
Check if each rack has an valid internal ring...................
[ERROR ]

Switches 0x21283a87cba0a0 0x21283a87b8a0a0 have 6 connections
between them
They should have at least 7 links between them
```

Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack

Recovery Appliance X3-2 systems and Recovery Appliance X2-2 systems have three Sun Datacenter InfiniBand Switch 36 switches. Starting with Recovery Appliance X4-2, Recovery Appliance Systems have two Sun Datacenter InfiniBand Switch 36 switches.

> **Note:**
>
> This procedure does not apply to Recovery Appliance X8M racks with RoCE Network Fabric.

The switch located in rack unit 1 (U1) is referred to as the spine switch. The other two switches are referred to as the leaf switches. The location of the leaf switches is as follows:

- Recovery Appliance Two-Socket Systems (X3-2 and later): rack unit 20 (U20) and rack unit 22 (U22)
- Recovery Appliance X2-2 racks: rack unit 20 (U20) and rack unit 24 (U24)
- Recovery Appliance Eight-Socket Systems (X2-8 and later) Full Racks: Rack unit 21 (U21) and rack unit 23 (U23)

The spine switch is the Subnet Manager Master for the InfiniBand Network Fabric subnet. The Subnet Manager Master has priority 8, and can be verified using the following procedure:

1. Log in to the spine switch as the `root` user.
2. Run the `setsmpriority list` command.

   The command should show that `smpriority` has a value of 8. If `smpriority` has a different value, then do the following:

   a. Use the `disablesm` command to stop the Subnet Manager.
   b. Use the `setsmpriority 8` command to set the priority to 8.
   c. Use the `enablesm` command to restart the Subnet Manager.

The leaf switches are the Standby Subnet Managers with a priority of 5. This can be verified using the preceding procedure, substituting a value of 5 in the `setsmpriority` command above.

> **Note:**
>
> Recovery Appliance Half Rack with Sun Fire X4170 Oracle Database Servers include two Sun Datacenter InfiniBand Switch 36 switches, which are set to priority 5.

To determine the Subnet Manager Master, log in as the `root` user on any InfiniBand Network Fabric switch, and run the `getmaster` command. The location of the Subnet Manager Master is displayed. The following is an example of the output from the `getmaster` command:

```
# getmaster
20100701 11:46:38 OpenSM Master on Switch : 0x0021283a8516a0a0 ports 36
Sun DCS 36
QDR switch dm01sw-ib1.example.com enhanced port 0 lid 1 lmc 0
```

The preceding output shows the proper configuration. The Subnet Master Manager is running on spine switch `dm01sw-ib1.example.com`.

If the spine switch is not the Subnet Manager Master, then do the following procedure to set the Subnet Manager Master:

1. Use the `getmaster` command to identify the current location of the Subnet Manager Master.
2. Log in as the `root` user on the leaf switch that is the Subnet Manager Master.
3. Disable Subnet Manager on the switch. The Subnet Manager Master relocates to another switch.
4. Use the `getmaster` command to identify the current location of the Subnet Manager Master. If the spine switch is not Subnet Manager Master, then repeat steps 2 and 3 until the spine switch is the Subnet Manager Master.
5. Enable Subnet Manager on the leaf switches that were disabled during this procedure.

Note:

- If the InfiniBand Network Fabric network consists of four or more racks cabled together, then only the spine switches should run Subnet Manager. The leaf switches should have Subnet Manager disabled on them.
- Recovery Appliance Half Racks with Sun Fire X4170 Oracle Database Servers, and Recovery Appliance Quarter Racks have two Sun Datacenter InfiniBand Switch 36 switches, and both are set to priority 5. The master is the one with the lowest GUID.

See Also:

- "Enable the Subnet Manager" in *Sun Datacenter InfiniBand Switch 36 User's Guide*
- "Disable the Subnet Manager" in *Sun Datacenter InfiniBand Switch 36 User's Guide*
- *Oracle Exadata Database Machine System Overview* for hardware component information
- Cabling tables in *Oracle Exadata Database Machine System Overview*

9

Installing the Recovery Appliance Software

This chapter describes the procedures for installing the Recovery Appliance software. It contains the following sections:

> **Note:**
>
> Recovery Appliance software is supported only when installed on Recovery Appliance hardware. See the *Zero Data Loss Recovery Appliance Licensing Information User Manual* for information about software licensing.

About the Installation Software

You use the following utilities to install the Recovery Appliance software:

- **`ra_preinstall.pl`**: Performs initial checks and prepares the Recovery Appliance for new installations and upgrades. This script was introduced in Zero Data Loss Recovery Appliance software update 12.1.1.1.7. See Zero Data Loss Recovery Appliance Upgrade and Patching (Doc ID 2028931.1) for more information.
- **`install.sh`**: Uses the files generated by Oracle Exadata Deployment Assistant to configure the networks, Oracle Database, and other software components.

  The OneCommand ZIP file that contains Oracle Exadata Deployment Assistant for use in Using the Deployment Assistant also contains this script. You can run this script from any system on the same network as Recovery Appliance. Some platforms use `install.cmd`.
- **`racli`**: Installs the Recovery Appliance software. You must run `racli` while logged in to a Recovery Appliance compute server, and not from another server on the same network.

> **Note:**
>
> Whenever the deployment configuration file is modified such as through a re-import and re-generation of XML, run `ra_preinstall.pl` again before running `install.sh`.

Software Installation Prerequisites

Before you begin the software installation, ensure that these procedures are complete:

- You used Oracle Exadata Deployment Assistant to generate the configuration files. See Using the Deployment Assistant .
- You assembled the installation files into the `WorkDir` directory under the OneCommand installation directory. (The OneCommand installation directory is the directory to which you extracted the OneCommand ZIP file, such as `linux-x64`.) See "#unique_189" for information about the OneCommand installation directory. For the list of installation files, see My Oracle Support Doc ID 1927416.1.
- The `checkip` script ran successfully with no remaining errors. See "#unique_190".
- All racks of a multirack installation are cabled together. See Extending Recovery Appliance with Multiple Racks.
- All components of all new racks are configured. See Configuring a Recovery Appliance Rack.
- Oracle Management Server is operating on the same network as Recovery Appliance.
- Reclaim the disk space on all compute servers. See "Reclaiming Disk Space on the Compute Servers".
- The Oracle tape library is installed and operational. A tape library is an optional component of Recovery Appliance.

About Adding a Rack to an Existing Rack

When you add a new Recovery Appliance rack to an older rack, the factory installed images are typically different. You can either update the older rack to the new image, or retain the older version and reimage the new rack to match it. In either case, ensure that all servers and all Sun Datacenter InfiniBand Switch 36 switches are at the same patch level.

About Adding New Storage Servers to an Existing Rack

Additional patching considerations include the Grid Infrastructure and database home releases and bundle patch updates. If the software installation applies new patches, then Oracle recommends changing the existing servers so the new servers inherit the releases as part of the extension procedure. This way, the number of servers that require patches is lower. Apply any patches to the existing servers in advance, so that they are at the desired level when the extension work is scheduled, thereby reducing the amount of work required by the extension.

Reclaiming Disk Space on the Compute Servers

Before you install the Recovery Appliance software, you must run the `reclaimdisks` script on each compute server. The script checks the volume configuration and reallocates all space correctly for the Recovery Appliance. Reclaiming the disk space converts the Linux system to four disks in a RAID-5 configuration at the disk controller level.

Perform the following steps to configure the disks on a compute server:

1. Log in to the compute server as the `root` user.
2. Change to the `/opt/oracle.SupportTools` directory.
3. Start the disk reclamation process:

   ```
   # ./reclaimdisks.sh -free -reclaim
   ```

   The command frees any preconfigured disks, schedules reclamation of all free disks to Linux on the next boot, and restarts Recovery Appliance. The process may take two hours to complete.

   > **Caution:**
   >
   > Do not restart Recovery Appliance or interrupt the reclamation process.

4. Validate the Linux configuration:

   ```
   ./reclaimdisks.sh -check
   ```

   The last lines of the output should be the following:

   ```
   [INFO]  Valid single boot configuration found for Linux: RAID5 from 4 disks
    with no global and dedicated hot spare disks
   Valid. Booted: Linux. Layout: Linux.
   ```

5. Repeat these steps on the other compute servers.

Installing the Software on Recovery Appliance

After completing the prerequisites and assembling the necessary files, you are ready to install the software.

To install the software on a new Recovery Appliance rack:

1. Log in as `root` to the first compute server.
2. Change to the `u01` directory:

   ```
   # cd /u01
   ```

3. Create a directory named `ra_install` directly under `u01`:

   ```
   # mkdir ra_install
   ```

4. Extract the latest Linux version of the OneCommand ZIP file into `/u01/ra_install/`, using the default directory name, `linux-x64`. (The ZIP file can be downloaded from My Oracle Support Doc ID 1927416.1.)
5. Copy the XML configuration files that were generated by using the Oracle Exadata Deployment Assistant into the `ExadataConfigurations` subdirectory under `/u01/ra_install/linux-x64`. See "#unique_189" for information about the configuration files.
6. Run the Recovery Appliance preinstallation script:
   a. Obtain the latest Recovery Appliance patch file from My Oracle Support Doc ID 1927416.1, and extract it into `/u01/ra_install/linux-x64/WorkDir`.
   b. Change to the `WorkDir` directory:

```
# cd /u01/ra_install/linux-x64/WorkDir
```

> **Notes:**
>
> - The preinstallation script is available with Zero Data Loss Recovery Appliance software update 12.1.1.1.7 and later. See Zero Data Loss Recovery Appliance Upgrade and Patching (Doc ID 2028931.1) for information about the latest patch file
> - The files you extracted to `WorkDir` should include a readme file (README.txt) that contains important information to review before you proceed

c. Run the script with the path of the XML configuration file and the Oracle Exadata Deployment Assistant directory:

```
# /usr/bin/perl ra_preinstall.pl --config_xml=/u01/ra_install/linux-x64/
ExadataConfigurations/Example_Inc-ra01.xml
--oeda_dir=/u01/ra_install/linux-x64
```

> **Note:**
>
> Before the `install.sh` has completed all steps, if the fix to a validation issue or any other issue requires re-generating the OEDA XML configuration file, run `ra_preinstall.pl` on the generated XML as given above before re-starting `install.sh` from the beginning.

7. Change to the `linux-x64` directory:

```
# cd /u01/ra_install/linux-x64
```

8. List the steps you need to run.

   This example runs the `install` script from a Linux system, using an XML file named `ExadataConfigurations/Example_Inc-ra01.xml`:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -l
```

9. Run each step in numeric order *except* the last step (Resecure Machine), and verify that it completed successfully before continuing to the next step. You will resecure the machine in a later step.

   This example runs step 1 from a Linux system, using an XML file named `ExadataConfigurations/Example_Inc-ra01.xml`:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -s 1
```

   See "About the Install Utility Steps".

10. Change to the Recovery Appliance bin directory:

```
# cd /opt/oracle.RecoveryAppliance/bin
```

11. Run the `racli` utility to install the Recovery Appliance software.

    This command runs all of the steps necessary to install the software based on the system and configuration settings:

```
# ./racli install appliance
```

Alternatively, you can run each step performed by this command individually by using the step option. However, ensure that you run all of the required steps in order and that each step completes successfully before you run the next step. The following example runs only the first installation step:

```
# ./racli install appliance --step=1
```

See "racli install appliance" for a description of each step.

> **Note:**
>
> If you encounter a problem while running the `install appliance` command, contact Oracle Support Services at `http://support.oracle.com` for assistance.

12. Change back to the `linux-x64` directory:

    ```
    # cd /u01/ra_install/linux-x64
    ```

13. Use the `install.sh` utility to run the last step (Resecure Machine) to secure Recovery Appliance.
14. To complete the Recovery Appliance installation:
    a. Change the default passwords for the Recovery Appliance database user (`rasys`) and the OSB tape backup users, if applicable.
    b. Use Oracle Enterprise Manager Cloud Control to deploy the agents and discover the Recovery Appliance targets. Then you can use Cloud Control to monitor and administer the Recovery Appliance environment.

       See "Setting Up Enterprise Manager to Administer Recovery Appliance".
    c. If you will use VLAN tagging for the Recovery Appliance ingest network, configure and test the VLAN tagging.

       See Enabling 8021.Q VLAN Tagging in Zero Data Loss Recovery Appliance over ingest networks (Doc ID 2047411.1) for instructions.
    d. Enroll the Oracle databases, so that they can use Recovery Appliance.

       See *Zero Data Loss Recovery Appliance Protected Database Configuration Guide*.

> **See Also:**
>
> - About Tape Backup Infrastructure
> - Changing OSB Tape Backup Passwords
> - Default User Accounts for Oracle Zero Data Loss Recovery Appliance

Install Utility Syntax

The `install` utility has the following syntax:

```
./install.sh -cf config_filename.xml option1 option2...
```

Install Utility Parameters

The syntax of the install utility enables you to run one step or multiple steps, and to undo one or more steps that ran with errors. You can then fix whatever problems caused the errors and run the step again.

Table 9-1 describes the options.

Table 9-1 Options of the Install Utility

Option	Description
`-cf` *`config_filename.xml`*	Identifies the XML configuration file. Required.
`-h`	Describes the script options.
`-l`	List the steps needed to install and configure the software for the configuration described in the XML file.
`-r` {*n*`-`*N* | *n*}	Runs steps *n* through *N*, or step *n*.
`-s` *n*	Runs step *n*.
-u {*n*`-`*N* | *n*}	Undoes steps *n* through *N*, or step *n*.

Install Utility Syntax Examples

The following examples provide examples of the syntax and sample output. The XML configuration file is named `Example_Inc-ra01.xml` (`linux-x64/ExadataConfigurations/Example_Inc-ra01.xml`).

This command displays Help for the utility:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -h
  install.sh -cf <config.xml> -l [options]
  install.sh -cf <config.xml> -s <step #> | -r <num-num>
  install.sh
  ARGUMENTS:
   -l                 List all the steps that exist
   -cf                config file name [Full path please]
   -s <step #>        Run the step # at a time
   -r <num-num>       Run the steps one after the other as long as no errors
                      are encountered
   -u <num-num> | <step#> Undo a range of steps or a particular step
   -h                 Usage
  Version : 14.254.04:00
```

This command lists the steps:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -l

1. Validate Configuration File
2. Setup Required Files
3. Create Users
4. Setup Cell Connectivity
5. Verify Infiniband
6. Calibrate Cells
7. Create Cell Disks
8. Create Grid Disks
```

```
9. Configure Alerting
10. Install Cluster Software
11. Initialize Cluster Software
12. Install Database Software
13. Relink Database with RDS
14. Create ASM Diskgroups
15. Create Databases
16. Apply Security Fixes
17. Install Exachk
18. Setup ASR Alerting
19. Create Installation Summary
20. Resecure Machine
```

The next example runs step 7:

```
./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -s 7

 Executing Create Cell Disks
 Check physical disks for errors before creating celldisks.............
 Restarting cell services..................
 Initializing cells.............
 Cleaning cells...................
 Restarting cell services...................
 Creating cell disks.......................
 Creating flashlog on cells..........
 Creating flashcache on cells........
 Successfully completed execution of step Create Cell Disks
]
```

This command reverses step 15:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -u 15

 Undoing Create Databases
 Deleting databases...
 Deleting database zdlra
      .
      .
      .
```

About the Install Utility Steps

The following descriptions show how each step contributes to the complete software installation. The install script lists the substeps as it performs them.

Step 1 Validate Configuration File

Runs a series of checks to ensure that the configuration file contains settings that are appropriate for the specified cluster on the current network:

- Validates the host names and IP addresses
- Verifies the operating system
- Validates the cluster networks and network connectivity
- Validates the network time protocol (NTP) setup
- Validates the physical disks in the storage servers
- Validates the compute servers for database readiness

Step 2 Setup Required Files

Ensures that the required files are accessible, copies them to the second compute server, unzips them as needed, and creates symbolic links.

See "Software Installation Prerequisites" if this step fails.

Step 3 Create Users

Creates the required users and groups for the cluster, updates `/etc/hosts`, and sets up a secure shell (SSH).

Step 4 Setup Cell Connectivity

Generates several files that are used by other steps, including `cellip.ora`, `cellinit.ora`, and `cellaffinity`, if required.

Step 5 Verify InfiniBand

Performs a series of checks on the InfiniBand fabric to ensure that the compute servers, storage servers, and InfiniBand switches are cabled together correctly.

If errors occur, a diagnostic ZIP file is generated. The output from this step includes the location of this file.

Step 6 Calibrate Cells

Uses the I/O calibration feature of Oracle Database to assess the I/O capability of the storage subsystem.

Step 7 Create Cell Disks

Performs a series of steps on the storage server disks that include creating the cell disks, and then creating a flash log and a flash cache on each cell.

Step 8 Create Grid Disks

Creates the grid disks for the cluster.

In this release, you must follow this step with additional manual procedures.

Step 9 Configure Alerting

Configures email on the storage servers, so that they can send alerts when problems occur.

Step 10 Install Cluster Software

Uses an Oracle Database utility to obtain the grid disk identifiers, and then runs the clusterware installer. It patches the software as needed.

Step 11 Initialize Cluster Software

If a replication network is being configured, this step creates the single client access name (SCAN) and virtual IP (VIP) for it. Then it initializes the cluster.

Step 12 Install Database Software

Installs the Oracle Database software and creates the Oracle Database home directory on both compute servers. It also patches the software as required.

In this release, you must follow this step with additional manual procedures.

Step 13 Relink Database with RDS

Relinks Oracle Database with Reliable Datagram Sockets (RDS).

Step 14 Create ASM Disk Groups

Obtains the grid disk identifiers and creates the Automatic Storage Management (ASM) disk groups.

Step 15 Create Databases

Creates the Recovery Appliance catalog on both compute servers.

Step 16 Apply Security Fixes

Restarts the cluster software to set the security parameters, and enables turbo mode if required.

Step 17 Install Exachk

Installs the `exachk` utility on the Recovery Appliance. See My Oracle Support Doc ID 1070954.1 for information about the utility.

Step 18 Setup ASR Alerting

Installs and configures Auto Service Request (ASR). Optional.

This step does the following:

- Installs the required software packages
- Configures the trap destinations
- Starts the monitoring daemon

Step 19 Create Installation Summary

Obtains the system details, calibrates the storage servers, and generates a summary XML report. The output from this step includes the location of this file.

Step 20 Resecure Machine

Replaces the default passwords with custom passwords.

Setting Up Enterprise Manager to Administer Recovery Appliance

To enable Enterprise Manager to administer Recovery Appliance, you must deploy the management agents to each compute server, then discover the targets for the Recovery

Appliance. The targets include the cluster, database, listeners, Automatic Storage Management (ASM), and the appliance itself.

Work with your Oracle support engineer to perform the tasks in this section.

> **Note:**
>
> See Doc ID 1929507.1 for information about the plug-ins and patches required for Zero Data Loss Recovery Appliance management, monitoring, and protected database management.

Deploying the Enterprise Manager Agents

Perform the following steps to deploy a management agent to each Recovery Appliance compute server:

1. On each compute server, create a directory owned by oracle:oinstall in which to install the agent.
2. Log in to Oracle Enterprise Manager Cloud Control (Cloud Control) as a user with privileges to discover and manage a Recovery Appliance target.
3. From the **Setup** menu, choose **Add Target**, select **Add Targets Manually**, choose **Add Host Targets**, then click **Add Host...**

   The Add Host Targets: Host and Platform page appears.
4. From the Platform list, select **Same for All Hosts**.
5. For each compute server, click **Add**, enter the fully qualified host name, and select **Linux x86-64** in the Platform list. Then click **Next**.

   The Add Host Targets: Installation Details page appears.
6. Specify the Agent installation directory that you created in Step 1, **Named Credential** for the user that will own the Agent installation, and other requested properties, then click **Next**.

   The Add Host Targets: Review page appears.
7. Click **Deploy Agent** to begin deploying Agents to all Recovery Appliance compute servers.

   The Add Host page appears, displaying warnings about remote prerequisite checks in the Agent Deployment Summary section. The Remote Prerequisite Check Details section indicates that the root.sh script could not be run due to insufficient privileges. You will fix this in a subsequent step.
8. To resume the deployment, click **Continue**, and select **All Hosts**.

   The Add Host page displays a message that the deployment was successful.
9. Follow to the on screen instructions to run root.sh manually on each compute server.
10. Return to the Add Host page, and click **Done**.

Discovering the Cluster

Perform the following steps to discover the Recovery Appliance cluster:

1. From the **Setup** menu in Cloud Control, choose **Add Target**, select **Add Targets Manually**.

   The Add Targets Manually page appears.

2. Select **Add Targets Using Guided Process**.

3. From the Target Types drop-down list, select **Oracle Cluster and High Availability Service**, then click **Add Using Guided Process...**

   The Add Target: Cluster and Oracle High Availability Service page appears.

4. Enter the host name of one of the compute servers on the Recovery Appliance, then click the Search icon.

   The Select Targets dialog appears.

5. Select the target host, and click **Select**.

6. Click **Discover Target**.

7. Review the automatically detected cluster information on the Cluster and Oracle High Availability Service page. Verify that all compute servers on the Recovery Appliance cluster are included in the list of cluster hosts and that the SCAN Name and SCAN Ports display the correct values for the ingest network, then click **Save**.

   When the cluster target is created successfully, the Confirmation window appears.

8. Close the Confirmation window.

Discovering the Cluster Database Targets

Before you discover the Recovery Appliance itself, perform the following steps to discover the Recovery Appliance cluster database, listener, and ASM targets:

1. From the **Setup** menu in Cloud Control, choose **Add Target**, select **Add Targets Manually**.

   The Add Targets Manually page appears.

2. Select **Add Targets Using Guided Process**.

3. From the Target Types drop-down list, choose **Oracle Database, Listener, and Automatic Storage Management**, then click **Add Using Guided Process...**

   The Database Discovery: Search Criteria page appears.

4. Use the Recovery Appliance cluster name or the host name of one of the compute servers as the criteria to perform the search, then click **Next**.

   The Database Discovery: Results page appears, showing the Recovery Appliance cluster database is listed, along with the ASM instances and listeners on all compute servers.

5. Select the cluster database, and click **Configure**.

6. Modify the settings to use fully qualified names in the Listener Machine Name fields and 1521 in the Port fields, then click **Save**.

7. Supply the monitoring credentials for the cluster database, and click **Test Connection**.

8. Select the cluster ASM, configure the appropriate monitoring credentials, and click **Test Connection**.
9. Ensure that the management database under Single Instance Databases (-MGMTDB) is not selected.
10. Select all listeners except the management listener (MGMTLSNR), and click **Next**.

    The Database Discovery: Review page appears.
11. Verify the information, then click **Save** to start monitoring the targets.
12. Close the Confirmation window.

Discovering the Recovery Appliance

After discovering the Recovery Appliance cluster and cluster database targets, the Recovery Appliance target itself can be discovered. As part of discovering the top-level Recovery Appliance target that will include all Enterprise Manager Recovery Appliance software and hardware management functionality, a separate Recovery Appliance hardware target will also be discovered.

Perform the following steps to manually discover the Recovery Appliance target:

1. From the **Setup** menu in Cloud Control, choose **Add Target**, then select **Add Targets Manually**.

   The Add Targets Manually page appears.
2. Select **Add Targets Using Guided Process**.
3. From the **Target Types** drop-down list, choose **Recovery Appliance**, then click **Add Using Guided Process...**

   The Recovery Appliance Hardware Discovery page appears.
4. Select **Discover new Recovery Appliance hardware components as targets**, and click **Discover Targets**.

   The Recovery Appliance Hardware Discovery: Discovery Inputs page appears.
5. Specify the URL of the agent on the compute server on which the schematic file (databasemachine.xml) resides. You can use the host name of the compute server on the Recovery Appliance to search for the correct value to use.

   Information about the schematic file appears.
6. Set the credential for the host, and click **Next**.

   The Recovery Appliance Hardware Discovery: Infiniband Discovery page appears.
7. Follow the detailed instructions in all subsequent pages of the Recovery Appliance Hardware Discovery wizard, supplying all requested credentials for hardware components. On the Review page, review all hardware component details, then click **Submit**. The Database Machine target representing the Recovery Appliance hardware will be created, along with targets for all hardware components.

   The Target Creation Summary page appears, showing a summary of all hardware targets created, including the Database Machine target representing the Recovery Appliance hardware and all member targets for all hardware components.
8. Click **Continue With Recovery Appliance Discovery**.

   The Recovery Appliance Discovery: Properties page appears.

9. Select the **Target Name** for the Recovery Appliance. The Recovery Appliance Hardware target name is already filled-in with the target name discovered in a previous step.

   (If the Recovery Appliance hardware discovery was completed without completing the full Recovery Appliance discovery process, that may leave one or more Recovery Appliance hardware targets in a state where they are not associated with a Recovery Appliance target. In this case, when Recovery Appliance discovery is re-initiated from the Add Targets Manually page, the discovery process will not automatically go in to the Recovery Appliance hardware discovery wizard. Instead, the process will go directly to the Properties page to allow selection of an unassociated Recovery Appliance hardware target. In this case, the hardware target name will not be pre-filled. Click the **Select Target** icon to launch the Search and Select Targets popup, which shows a list of Recovery Appliance hardware targets that have not yet been associated with a Recovery Appliance target. Select the correct **Recovery Appliance hardware target**.)

10. In the Recovery Appliance Administrator Credentials section, specify the database user credentials for the Recovery Appliance recovery catalog owner. This user is also the Recovery Appliance administrator. You can use **Named** or **New** credentials. Click **More Details** to see detailed information about the credentials.

11. In the Recovery Appliance Monitoring Credentials section, specify the database user credentials that will be used to monitor the Recovery Appliance. If necessary, that user will be granted the role required to access monitoring information in the Recovery Appliance database. Alternatively, you can select **Use base catalog user credentials as monitoring credentials** which disables the Username and Password fields.

12. Click **Next**.

    The Recovery Appliance Discovery: Oracle Secure Backup Domain page appears.

13. If Oracle Secure Backup is installed on the Recovery Appliance, enter /usr/local/oracle/backup in the Installation Home field and specify the monitoring credentials for the Oracle Secure Backup domain. Otherwise, select **Skip Oracle Secure Backup Domain Discovery**.

14. Click **Next**.

    The Recovery Appliance Discovery: Review page appears.

15. Review the target discovery selections, and click **Submit**.

    The Recovery Appliance target is created.

Part III

Security and Maintenance of Recovery Appliance

Part III provides information about security and maintenance procedures for Recovery Appliance. It contains the following chapters:

- User Security on Recovery Appliance
- Keeping the Recovery Appliance Secure
- Maintaining the Recovery Appliance Hardware
- Reconfiguring the Hardware

10

User Security on Recovery Appliance

Increase the security of your data and system by limiting user access and developing strong password security policies.

Default User Accounts for Oracle Zero Data Loss Recovery Appliance

The following table lists the default users and passwords for the Oracle Zero Data Loss Recovery Appliance components. All default passwords should be changed after installation of the Recovery Appliance.

Table 10-1 Default Users and Passwords for Oracle Zero Data Loss Recovery Appliance

Component	User Name and Password
Compute servers	Operating system users: • `root/welcome1` • `oracle/We1come$` • `dbmadmin/welcome` • `dbmmonitor/welcome` • `raext/(locked and blocked from SSH access)` • `railm/(locked and blocked from SSH access)` • Password for the GRUB boot loader: `sos1Exadata` Database users: • `SYS/We1come$` • `SYSTEM/We1come$` • `raext/(externally authenticated)` • `ralim/(externally authenticated)` • `rasys/change^Me2` OSB tape backup application users: • `admin/welcome1` • `oracle/welcome1` • encryption key wallet/`welcome1`

Table 10-1 (Cont.) Default Users and Passwords for Oracle Zero Data Loss Recovery Appliance

Component	User Name and Password
Storage servers	• `root/welcome1` • `celladmin/welcome` • `cellmonitor/welcome` • `CELLDIAG` `CELLDIAG` is an Exadata storage software user, not an operating system user. The password of the `CELLDIAG` user is reset to a random password during the "Apply Security Fixes" step of Oracle Exadata Deployment Assistant. If this step is not run, then the default password is `Welcome12345`. • Password for the GRUB boot loader: `sos1Exadata`
RoCE Network Fabric	• `root/welcome1`
InfiniBand Network Fabric switches	• `root/welcome1` • `nm2user/changeme` • `ilom-admin/ilom-admin` • `ilom-operator/ilom-operator`
Ethernet switches	`admin/welcome1` **Note:** Secure the `enable mode` *`password`* and *`secret`* values for the `admin` user.
Power distribution units (PDUs)	• `admin/welcome1` The password for the `admin` user is `adm1n` if you reset the PDU to factory default settings.
Compute server ILOMs	• `root/welcome1` • `MSUser` Management Server (MS) uses this account to manage ILOM and reset it if it detects a hang. Do not modify this account. This account is to be used by MS only. Each time MS starts up, it deletes the previous `MSUser` account and re-creates the account with a randomly generated password. The `MSUser` password is not persisted anywhere. If you need to change account passwords regularly, you can restart MS to change the password of the `MSUser` account.
Storage server ILOMs	• `root/welcome1` • `MSUser` See the description above for details about this user.
InfiniBand Network Fabric ILOMs	• `ilom-admin/ilom-admin` • `ilom-operator/ilom-operator` • `root/welcome1`

> **Note:**
>
> After the Recovery Appliance has been deployed, the installation process disables all root SSH keys and expires all user passwords as a security measure for your system. If you do not want the SSH keys disabled or the passwords expired, advise the installation engineer before the deployment.

> **See Also:**
>
> "Changing Component Passwords" to learn how to change the passwords for the Recovery Appliance components.

Default Password Requirements

Oracle Exadata Deployment Assistant (OEDA) implements a default password policy on Oracle Exadata Database Machine.

The last step of OEDA, "Secure Oracle Exadata Database Machine", implements the following password requirements:

- Dictionary words are not valid or accepted.
- Character classes for passwords are uppercase letters, lowercase letters, digits, and special characters.
- Passwords must contain characters from all four character classes. Passwords using only one, two, or three character classes are not allowed.
- The minimum length of a password is eight characters.
- Pass-phrases are allowed. A pass-phrase should contain at least three words, be 16 to 40 characters in length, and contain different character classes.
- A new password cannot be similar to old passwords. There must be at least eight characters in the new password that were not present in the old password.
- A maximum of three consecutive characters of the same value can be used in a password.
- A maximum of four consecutive characters of the same character class can be used in a password. For example, `abcde1#6B` cannot be used as a password because it uses five consecutive lower case letters.

Default Security Settings Enacted by OEDA

Oracle Exadata Deployment Assistant (OEDA) includes a step to increase hardware security on Recovery Appliance.

The last step of OEDA implements the following security policies:

- For all newly created operating system users on the compute servers and storage servers, the following password-aging values are set:
  - The maximum number of days for a password is 60 days.
  - The minimum amount of time between password changes is 24 hours.

  - The number of days of alerts before a password change is seven days.
  - All non-root users must change their password at their next log in.
- An operating system user account is temporarily locked for 10 minutes after one failed log in attempt.
- An operating system user account is locked after five failed attempts.
- For the `root` user, SSH equivalency is removed for all compute servers and storage servers.
- The following permissions are set by OEDA:
  - The Automatic Diagnostic Repository (ADR) base directory, `$ADR_BASE`, has SUID (Set owner User ID) on the `diag` directory and its sub-directories.
  - The `celladmin` user group has read and write permissions on the `$ADR_BASE`.

11
Keeping the Recovery Appliance Secure

This chapter describes policies and procedures to keep Recovery Appliance secure.

Securing the Hardware

After installation of Oracle Exadata Database Machine, the hardware should be secured.

Hardware can be secured by restricting access to the hardware and recording the serial numbers. Oracle recommends the following practices to restrict access:

- Install Oracle Exadata Database Machine and related equipment in a locked, restricted-access room.
- Lock the rack door unless service is required on components within the rack.
- Restrict access to hot-pluggable or hot-swappable devices because the components can be easily removed by design. See
- Store spare field-replaceable units (FRUs) or customer-replaceable units (CRUs) in a locked cabinet. Restrict access to the locked cabinet to authorized personnel.
- Mark all significant items of computer hardware, such as FRUs.
- Keep hardware activation keys and licenses in a secure location that is easily accessible to the system managers in the case of a system emergency.
- Record the serial numbers of the components in Oracle Exadata Database Machine, and keep a record in a secure place. All components in Oracle Exadata Database Machine have a serial number.

Related Topics

- How To Obtain The Serial Number Associated With The System Board, Motherboard, Disk Controller, Disks, Infiniband HCA And More Contained In A Cell Or Compute Box (Exadata-Sun V2 or X2 / 11.2)? (My Oracle Support Doc ID 949614.1)
- How to Determine the Serial Number of a Datacenter InfiniBand Switch 36 or QDR InfiniBand Gateway InfiniBand Switch (My Oracle Support Doc ID 1299791.1)

Getting the Rack Serial Number

Use the ipmitool utility to get the serial number for the rack.

When interacting with Oracle Support Services, the CSI number for a rack is based on the rack serial number.

1. Log in to one of the servers in the rack as the `root` user.
2. Use `ipmitool` to get the serial number for the rack.

```
# ipmitool sunoem cli "show /SP system_identifier"
Connected. Use ^D to exit.
-> show /SP system_identifier
```

```
 /SP
    Properties:
        system_identifier = Exadata Database Machine X2-8xxxxAKyyyy


-> Session closed
Disconnected
```

Getting the Serial Numbers for Rack Components

The CheckHWnFWProfile command can be used to display the serial number of most of the system components.

1. Log in to one of the servers in the rack as the `root` user.
2. On each server in the rack, use `CheckHWnFWProfile` with the `-S` option to display the serial number of the components for that server.

   ```
   # /opt/oracle.SupportTools/CheckHWnFWProfile -S > /tmp/
   CheckHWnFWProfile_hostname.txt
   ```

   The result is specific to each server, so the command must be performed on every node. The following is a partial example of the output:

   ```
   Server_Model=ORACLE_SERVER_X8-2L
   ====START SERIAL NUMBERS====
   ==Motherboard, from dmidecode==
   --System serial--
   1904XCA000
   --Motherboard serial--
   469996N+0000RD01RN
   --Chassis serial--
   1900XCA000
   --Rack serial--
   AK00400000
   ==Infiniband HCA==
   ID:      CX354A - ConnectX-3 QSFP
   PN:      7046442
   EC:      XX
   SN:      465000K-1800000000
   V0:      PCIe Gen3 x8
   ==Motherboard, RAM etc from ipmitool==
   FRU Device Description : Builtin FRU Device (LUN 0 ID 0)
   ...
    Product Name          : ILOM
    Product Version       : 4.0.4.38.a

   FRU Device Description : BMC
   ...
    Product Name          : ILOM
     Product Version       : 4.0.4.38.a

   FRU Device Description : /SYS (LUN 0 ID 3)
   ```

```
...
 Product Part Number   : 8200669
 Product Serial        : 1900XCA000

FRU Device Description : DBP (LUN 0 ID 210)

 Board Part Number     : 7341141
 Board Extra           : Rev 09

FRU Device Description : HDD0 (LUN 0 ID 47)
 Device not present (Requested sensor, data, or record not found)

FRU Device Description : HDD1 (LUN 0 ID 48)
 Device not present (Requested sensor, data, or record not found)

...

FRU Device Description : MB (LUN 0 ID 4)
 Board Mfg Date        : Sun Jan 20 16:57:00 2019
 Board Mfg             : Oracle Corporation
...

FRU Device Description : MB/BIOS (LUN 0 ID 5)
...

FRU Device Description : MB/CPLD (LUN 0 ID 8)
 Product Manufacturer  : Oracle Corporation
 Product Name          : Power Control FPGA
 Product Version       : FW:3.9

FRU Device Description : M2R0/SSD0 (LUN 0 ID 211)
 Device not present (Requested sensor, data, or record not found)

FRU Device Description : M2R1/SSD0 (LUN 0 ID 212)
 Device not present (Requested sensor, data, or record not found)

FRU Device Description : MB/NET0 (LUN 0 ID 43)
 Product Manufacturer  : INTEL
 Product Name          : 1G Ethernet Controller
...

FRU Device Description : MB/P0 (LUN 0 ID 16)
 Product Manufacturer  : Intel
 Product Name          : Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
...

FRU Device Description : MB/P0/D0 (LUN 0 ID 24)
 Product Manufacturer  : Samsung
 Product Name          : 16384MB DDR4 SDRAM DIMM
...

FRU Device Description : MB/P0/D1 (LUN 0 ID 25)
 Device not present (Requested sensor, data, or record not found)

FRU Device Description : MB/P0/D2 (LUN 0 ID 26)
```

```
 Product Manufacturer  : Samsung
 Product Name          : 16384MB DDR4 SDRAM DIMM
...


FRU Device Description : MB/P1 (LUN 0 ID 17)
 Product Manufacturer  : Intel
 Product Name          : Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
...

FRU Device Description : MB/P1/D0 (LUN 0 ID 36)
 Product Manufacturer  : Samsung
 Product Name          : 16384MB DDR4 SDRAM DIMM
...
FRU Device Description : PS0 (LUN 0 ID 63)
...
FRU Device Description : PS1 (LUN 0 ID 64)
...
FRU Device Description : SP/NET0 (LUN 0 ID 1)
...
FRU Device Description : SP/NET1 (LUN 0 ID 2)
...
FRU Device Description : /UUID (LUN 0 ID 6)
...
FRU Device Description : TOP_LEVEL_CH (LUN 0 ID 251)
 Chassis Type          : Rack Mount Chassis
 Chassis Part Number   : 8200669
 Chassis Serial        : 1900XCA0000
 Chassis Extra         : chassis_name:ORACLE SERVER X8-2L

FRU Device Description : TOP_LEVEL_PROD (LUN 0 ID 250)
 Product Manufacturer  : Oracle Corporation
 Product Name          : Exadata X8-2
 Product Part Number   : Exadata X8-2
 Product Serial        : AK00430000

====END SERIAL NUMBERS====
```

Getting the Rack Serial Number for a Cisco 9336C or 9348 Switch

Use the `show license host-id` command on the switch to get the serial number.

1. Connect to the switch from a server with SSH equivalency configured, or log in as the `admin` user.
2. Obtain the serial number for the switch by entering the `show license host-id` command.

   The host ID is also referred to as the device serial number.

   ```
   # switch# show license host-id
   License hostid: VDH=FLA12345678
   ```

   Use the entire ID that appears after the equal sign (=). In this example, the host ID is FLA12345678.

Getting the Rack Serial Number for a Sun Datacenter InfiniBand Switch 36

Use the `showfruinfo` command on the switch to get the serial number.

1. Log in to the switch as `root`.

   ```
   $ ssh root@switch_name
   ```

2. Use the `showfruinfo` command to view the serial number for the switch.

   ```
   root@ib-switch-> showfruinfo
   Sun_Man1R:
   UNIX_Timestamp32 : Fri Mar 19 16:29:59 2010
   Sun_Fru_Description : ASSY,NM2-GW
   Vendor_ID_Code : 11 E1
   Vendor_ID_Code_Source : 01
   Vendor_Name_And_Site_Location : 4577 CELESTICA CORP. SAN JOSE CA US
   Sun_Part_Number : 5111402
   Sun_Serial_Number : 0110SJC-1010NG0040
   Serial_Number_Format : 4V3F1-2Y2W2X4S
   Initial_HW_Dash_Level : 03
   Initial_HW_Rev_Level : 50
   Sun_Fru_Shortname : NM2 gateway
   Sun_Hazard_Class_Code : Y
   Sun_SpecPartNo : 885-1655-01
   Sun_FRU_LabelR:
   Sun_Serial_Number : AK000XXXX2
   FRU_Part_Dash_Number : 541-4188-01
   ```

Getting the Serial Number for a Cisco 4948 Ethernet Switch

Use the `sh inventory` command on the switch to get the serial number.

1. Log in to the Cisco Ethernet switch.
2. Obtain the serial number for the switch and its components by entering the `sh inventory` command.

   ```
   # Switch# sh inventory
   NAME: "Switch System", DESCR: "Cisco Systems, Inc. WS-C4948 1 slot switch
   "
   PID:                    , VID:      , SN: FOX0000G0B6
   NAME:  "Linecard(slot 1)", DESCR: "10/100/1000BaseT (RJ45), 1000BaseX
   (SFP)
    Supervisor with 48 10/100/1000BASE-T ports and 4 1000BASE-"
   PID: WS-C4948           , VID: V09  , SN: FOX0000G0B6
   NAME: "Power Supply 1", DESCR: "Power Supply ( AC 300W )"
   PID: PWR-C49-300AC      , VID:      , SN: QCS0000B1XR
   NAME: "Power Supply 2", DESCR: "Power Supply ( AC 300W )"
   PID: PWR-C49-300AC      , VID:      , SN: QCS0000B1X5
   ```

Securing the Software

Frequently, hardware security is implemented through software measures.

Implement the following guidelines to protect hardware and software:

- Change all default passwords when the system is installed at the site. Recovery Appliance uses default passwords for initial installation and deployment that are widely known. A default password could allow unauthorized access to the equipment. Devices such as the network switches have multiple user accounts. Be sure to change all account passwords on the components in the rack.
- Limit use of the `root` super user account. Use `non-root` access when possible. Create and use Integrated Lights Out Manager (ILOM) user accounts for individual users to ensure a positive identification in audit trails, and less maintenance when administrators leave the team or company.
- Restrict physical access to USB ports, network ports, and system consoles. Servers and network switches have ports and console connections, which provide direct access to the system.
- Restrict the capability to restart the system over the network.

Maintaining a Secure Environment

After security measures are implemented, they must be maintained to keep the system secure.

Software, hardware and user access need to be updated and reviewed periodically. For example, organizations should review the users and administrators with access to Recovery Appliance to verify if the levels of access and privilege are appropriate. Without review, the level of access granted to individuals may increase unintentionally due to role changes or changes to default settings. It is recommended that access rights for operational and administrative tasks be reviewed to ensure that each user's level of access is aligned to their roles and responsibilities.

Refer to User Accounts in the Recovery Appliance Environment.

Organizations are encouraged to utilize tools to detect unauthorized changes, configuration drift, and prepare for security updates. Oracle Enterprise Manager provides an integrated solution for managing operational issues for hardware, deployed applications, and services.

Maintaining Network Security

After the networks are configured based on the security guidelines, regular review and maintenance is needed.

The management network switch configuration file should be managed offline, and access to the configuration file should be limited to authorized administrators. The configuration file should contain descriptive comments for each setting. Consider keeping a static copy of the configuration file in a source code control system. Periodic reviews of the client access network are required to ensure that secure host and Integrated Lights Out Manager (ILOM) settings remain intact and in effect. In addition, periodic reviews of the settings ensure that they remain intact and in effect.

Follow these guidelines to ensure the security of local and remote access to the system:

- Set time-outs for extended sessions and set privilege levels.
- Use authentication, authorization, and accounting (AAA) features for local and remote access to a switch.
- Use the port mirroring capability of the switch for intrusion detection system (IDS) access.
- Implement port security to limit access based upon a MAC address.
- Require users to use strong passwords by setting minimum password complexity rules and password expiration policies.
- Enable logging and send logs to a dedicated secure log host.
- Configure logging to include accurate time information, using NTP and timestamps.
- Review logs for possible incidents and archive them in accordance with the organization's security policy.

Standard 140 of FIPS (Federal Information Processing Standards) relates to security and cryptography. FIPS 140 is a collection of standards published by NIST (National Institute of Standards and Technology), an agency of the United States federal government. FIPS 140 protects data during transit as well as at rest. It specifies security standards for cryptographic components within a computing environment. FIPS 140 is useful for organizations that need to document that their computing environment meets a published level of security. Many government agencies and financial institutions use FIPS 140 qualified systems.

Configuring FIPS 140 at the Oracle Database level enables the use of FIPS 140 cryptographic modules in the Secure Sockets Layer (SSL), transparent data encryption (TDE), DBMS_CRYPTO PL/SQL package, and Exadata Smart Scan. This protects data while processing Smart Scan offload operations.

12

Maintaining the Recovery Appliance Hardware

This chapter describes how to maintain the Recovery Appliance rack components. It contains the following topics:

> **See Also:**
>
> Replacement Units

Cautions and Warnings

When maintaining the Recovery Appliance hardware, observe the following precautions:

> **WARNING:**
>
> Do not touch the parts of this product that use high-voltage power. Touching them might result in serious injury.

Caution:

- Do not power off Recovery Appliance unless there is an emergency. In that case, follow "Emergency Power-Off Procedure".
- Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to the hardware components.
- Keep the top, front, and back of the cabinets clear to allow proper airflow and to prevent the components from overheating.

Determining the Server Model

Use the following command to determine the model of a compute server or a storage server:

```
dmidecode -s system-product-name
```

Powering On and Off a Recovery Appliance Rack

This section includes the following topics:

- Emergency Power-Off Procedure
- Shutting Down Recovery Appliance
- Starting Up Recovery Appliance

Emergency Power-Off Procedure

In an emergency, halt power to Recovery Appliance immediately. The following emergencies might require powering off Recovery Appliance:

- Natural disasters, such as earthquake, flood, hurricane, tornado, or cyclone
- Abnormal noise, smell, or smoke coming from the system
- Threat to human safety

Powering Off in an Emergency

In an emergency, do one of the following:

- Turn off power at the circuit breaker.
- Pull the emergency power-off switch in the computer room.

After the emergency, contact Oracle Support Services about restoring power to the system.

About the Emergency Power-Off Switch

You can use the emergency power-off (EPO) switch to remove power from Recovery Appliance.

EPO switches are required when computer equipment contains batteries capable of supplying more than 750 volt-amperes for more than five minutes. Systems that have these batteries include internal EPO hardware for connecting to a site EPO switch or relay.

Shutting Down Recovery Appliance

Under normal, nonemergency conditions, you can power down the software services and hardware gracefully.

Stop all software services before shutting down the rack components.

Stopping Recovery Appliance Services

You must stop the Recovery Appliance services, Oracle Database File System, Oracle Database, and the cluster services.

To stop the Recovery Appliance services:

1. Disable the keystore as part of appliance shutdown. If the keystore is not created, skip this step.

   If a keystore is created, this is required to stop the HSM wallet for copy to cloud encryption.

   Before appliance shutdown, run:

   ```
   [root@myhost ~]# racli disable keystore
   [root@myhost ~]# racli status keystore
          Node: zdlra41
   Wallet Type: HSM
        Status: Closed
          Node: zdlra42
   Wallet Type: HSM
        Status: Closed
   ```

2. Log in as `oracle` to either Recovery Appliance compute server.
3. Open a SQL connection to Oracle Database as the `rasys` user:

   ```
   $ sqlplus rasys
   ```

4. Check the status of the services:

   ```
   SQL> SELECT state FROM ra_server;

   STATE
   ------------
   ON
   ```

5. Shut down Recovery Appliance services:

   ```
   SQL> exec dbms_ra.shutdown;
   ```

6. Disconnect from Oracle Database:

```
SQL> exit
```

7. If Oracle Secure Backup is configured:
   a. Switch to the `root` user.
   b. Check the current status of Oracle Secure Backup:

```
# $GRID_HOME/bin/crsctl status res osbadmin
  NAME=osbadmin
  TYPE=cluster_resource
  TARGET=ONLINE
  STATE=ONLINE on example01adm04
```

   c. If Oracle Secure Backup is online, then stop it:

```
# $GRID_HOME/bin/crsctl stop res osbadmin
```

   d. Switch back to the `oracle` user.
8. Check the status of Oracle Database:

```
$ srvctl status database -d zdlra5
Instance zdlra51 is running on node radb07
Instance zdlra52 is running on node radb08
```

9. Stop Oracle Database:

```
$ srvctl stop database -d zdlra5
```

10. Verify that Oracle Database is stopped:

```
$ srvctl status database -d zdlra5
Instance zdlra51 is not running on node radb07
Instance zdlra52 is not running on node radb08
```

11. Switch to the `root` user.
12. Stop the Oracle Clusterware stack on all nodes in the cluster:

```
# $GRID_HOME/bin/crsctl stop cluster -all
CRS-2673: Attempting to stop 'ora.crsd' on 'zdlradb07'
CRS-2790: Starting shutdown of Cluster Ready Services-managed
resources on
'zdlradb07'
CRS-2673: Attempting to stop 'ora.LISTENER_SCAN2.lsnr' on
'zdlradb07'
CRS-2673: Attempting to stop 'ora.LISTENER_SCAN1.lsnr' on
'zdlradb07'
     .
     .
     .
#
```

If the command fails, reenter it with the `-f` option.

13. On each compute server, run the following command to stop the Oracle Cluster Ready Services (CRS):

```
# $GRID_HOME/bin/crsctl stop crs
CRS-2791: Starting shutdown of Oracle High Availability Services-managed
resources on 'radb08'
CRS-2673: Attempting to stop 'ora.crf' on 'radb08'
CRS-2673: Attempting to stop 'ora.mdnsd' on 'radb08'
     .
     .
     .
CRS-2677: Stop of 'ora.crf' on 'radb08' succeeded
CRS-2677: Stop of 'ora.mdnsd' on 'radb08' succeeded
CRS-2793: Shutdown of Oracle High Availability Services-managed resources
on 'radb08' has completed
CRS-4133: Oracle High Availability Services has been stopped.
```

14. Shut down or reboot the hardware as required, in the following order:
    a. Compute servers
    b. Storage servers
    c. Rack and switches

Powering Down the Servers

Before powering down a server, stop the services running on it, as described in "Shutting Down Recovery Appliance".

To shut down a compute server or a storage server:

1. Log in to the server as `root`.
2. Stop the operating system:

```
# shutdown -h -y now
```

   Or restart the operating system:

```
# shutdown -r -y now
```

Example 12-1 Powering Off Recovery Appliance Using the dcli Utility

1. Stop Oracle Clusterware on all compute servers:

```
# GRID_HOME/grid/bin/crsctl stop cluster -all
```

2. Shut down the other compute server in the rack:

```
# dcli -l root -g ra-adm02 shutdown -h -y now
```

   In the preceding command, `ra01adm02` is the name of the second compute server.

3. Shut down all storage servers:

```
# dcli -l root -g cell_group shutdown -h -y now
```

   In the preceding command, `cell_group` is a file that lists all storage servers.

4. Shut down the local compute server:

```
shutdown -h -y now
```

5. Power off the rack.

Use the `dcli` utility to run the `shutdown` command on multiple servers simultaneously. Do not run `dcli` from a server that will be powered off by the command.

The following example shuts down a group of storage servers listed in a file named `cell_group`:

```
# dcli -l root -g cell_group shutdown -h -y now
```

Example 12-1 shows the power off procedure for the rack when using the `dcli` utility to shut down multiple servers simultaneously. The commands run from a compute server.

Powering the Network Switches

The gateway and spine switches do not have power controls. They power off when power is removed, by turning off a PDU or a breaker in the data center.

Starting Up Recovery Appliance

Turn on the rack components first, then start the software services.

Starting Up Recovery Appliance Components

To power on the rack components, use one of the following methods:

- Press the power button on the front of the component.
- Log in to Oracle ILOM and apply power to the system. See "Powering On Servers Remotely".

Startup Sequence

Power on the rack components in this sequence:

1. Rack and switches

   Allow the switches a few minutes to initialize, before you start the storage servers.

2. Storage servers

   Allow five to 10 minutes for the storage servers to start all services. Ensure that they finish initializing before you start the compute servers.

3. Compute servers

   When a compute server is powered on, the operating system and Oracle Clusterware start automatically. Oracle Clusterware then starts all resources that are configured to start automatically.

Powering On Servers Remotely

You can use the Oracle ILOM interface to power on the Recovery Appliance servers remotely. To access Oracle ILOM, use the web console, the command-line interface (CLI), intelligent platform management interface (IPMI), or simple network management protocol (SNMP).

For example, to apply power to server `ra01cel01` using IPMI, you use its Oracle ILOM with a command like the following:

```
# ipmitool -H ra01cel01-ilom -U root chassis power on
```

IPMItool must be installed on the server where you use the command.

> **See Also:**
>
> Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation for additional information about using Oracle ILOM to power on the servers:
>
> http://docs.oracle.com/cd/E19860-01/index.html

Starting the Recovery Appliance Software

1. Log in as `root` to a Recovery Appliance compute server.
2. Confirm that Oracle Cluster Ready Services (CRS) is running:

```
# $GRID_HOME/bin/crsctl status server
  NAME=radb07
  STATE=ONLINE

  NAME=radb08
  STATE=ONLINE
```

3. If CRS is not running, then start it:

```
# $GRID_HOME/bin/crsctl start cluster -all
CRS-2672: Attempting to start 'ora.evmd' on 'radb07'
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'radb07'
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'radb08'
     .
     .
     .
#
```

4. Switch to the `oracle` user.
5. Verify that Oracle Database is running:

```
$ srvctl status database -d zdlra5
Instance zdlra51 is not running on node radb07
Instance zdlra52 is not running on node radb08
```

6. If Oracle Database is not running:
   a. Start Oracle Database:

```
$ srvctl start database -d zdlra5
```

b. Confirm that Oracle Database is running:

```
$ srvctl status database -d zdlra5
Instance zdlra51 is running on node radb07
Instance zdlra52 is running on node radb08
```

7. If Oracle Secure Backup is enabled, then start it:

```
# $GRID_HOME/bin/crsctl start res osbadmin
```

8. Connect to Oracle Database as the RASYS user:

```
$ sqlplus rasys
```

9. Check the status of Recovery Appliance services:

```
SQL> SELECT state FROM ra_server;

STATE
------------
OFF
```

10. If the services are off, then start them:

```
SQL> exec dbms_ra.startup;
```

11. Confirm that the services are started:

```
SQL> /

STATE
------------
ON
```

12. Enable the keystore as part of appliance startup. This is required step to open the HSM wallet for copy to cloud encryption. After appliance restart, run:

```
[root@myhost ~]# racli enable keystore
[root@myhost ~]# racli status keystore
       Node: zdlra42
Wallet Type: HSM
     Status: Open
       Node: zdlra41
Wallet Type: HSM
     Status: Open
```

Replacing the Disk Controller Batteries

The disk controllers in storage servers and compute servers have battery-backed write cache to accelerate write performance. If the battery charge capacity degrades, so that the battery cannot protect the cached data for a power loss of 48 hours or more, then the write cache is disabled and the disk controller switches to write-through mode. Write performance is reduced, but no data is lost.

The battery charge capacity degrades over time, and its life expectancy is inversely proportional to the operating temperature. Table 12-1 shows the worst case life expectancy of a battery in Recovery Appliance.

Table 12-1 Battery Life Expectancy

Inlet Ambient Temperature	Battery Lifetime
< 25 degrees Celsius (77 degrees Fahrenheit)	3 years
< 32 degrees Celsius (89.6 digresses Fahrenheit)	2 years

Monitoring Batteries in the Compute Servers

To monitor the battery change capacity in the compute servers:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -a0 | grep "Full Charge" -A5 | sort \
| grep Full -A1
```

The following is an example of the output from the command:

```
Full Charge Capacity: 1357 mAh
Max Error: 2 %
```

You should proactively replace batteries that have a capacity less than 800 milliampere hour (mAh) and a maximum error less than 10%. Immediately replace any battery that has less than 674 mAh or a maximum error greater than 10%.

To monitor the battery temperature:

```
/opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -a0 | grep BatteryType; \
/opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -a0 | grep -i temper
```

The following is an example of the output from the command:

```
BatteryType: iBBU08
Temperature: 38 C
  Temperature                 : OK
  Over Temperature        : No
```

If the battery temperature is greater than or equal to 55 degrees Celsius, then determine the cause, and correct the problem.

> **Note:**
>
> Storage servers generate an alert when the battery charge capacity is insufficient, the temperature is high, or the battery should be replaced.

Replacing Batteries in Disk Controllers

Oracle replaces the failed batteries at no extra charge under these conditions:

- The battery charge capacity in the disk controllers falls below the minimum threshold
- The system is covered either by the Oracle Premier Support for Systems or occurs during the warranty period.

For customers with Premier Support for Systems, Oracle attempts to proactively replace the batteries in Recovery Appliance before the end of the estimated lifetime, on a best effort basis.

Replacing a Power Distribution Unit

A power distribution unit (PDU) can be replaced while Recovery Appliance is online. The second PDU in the rack maintains the power to all components in the rack. PDU-A is on the left, and PDU-B is on the right, when viewing the rack from the rear.

PDU Replacement Guidelines

Before replacing a PDU, review the following guidelines to ensure that you can perform the procedure safely and without disrupting availability:

- Unlatching the InfiniBand cables while removing or inserting PDU-A might remove servers from the cluster and thus make the rack unavailable. Be careful when handling the InfiniBand cables, which are normally latched securely. Do not place excessive tension on the InfiniBand cables by pulling them.
- Unhooking the wrong power feeds shuts down the rack. Trace the power cables that will be replaced from the PDU to the power source, and only unplug those feeds.
- Allow time to unpack and repack the PDU replacement parts. Notice how the power cords are coiled in the packaging, so you can repack the failed unit the same way.
- Removing the side panel decreases the time needed to replace the PDU. However, removing the side panel is optional.
- Using a cordless drill or power screwdriver decreases the time needed to replace the PDU. Allow more time for the replacement if you use a hand wrench. A screwdriver requires Torx T30 and T25 bits.
- You might need to remove the server cable arms to move the power cables. In that case, twist the plug connection and flex the cable arm connector, to avoid having to unclip the cable arm. If you must unclip the cable arm, then support the cables with one hand, remove the power cord, and then clip the cable arm. Do not leave the cable arm hanging.
- When you remove the T30 screws from the L-bracket, do not remove the T25 screws or nuts that attach the PDU to the bracket, until the PDU is out of the rack.

Replacing a PDU

To replace a PDU:

1. Restart the PDU monitor to identify the network settings:
   a. Press the reset button for 20 seconds, until it starts to count down from 5 to 0. While it is counting, release the button, and then press it once.
   b. Record the network settings, firmware version, and so on, displayed on the LCD screen as the monitor restarts.

   If the PDU monitor is not working, then retrieve the network settings by connecting to the PDU over the network, or from the network administrator.

2. Turn off all PDU breakers.
3. Unplug the PDU power plugs from the AC outlets.

   If the rack is on a raised floor, then move the power cords out through the floor cutout. You might need to maneuver the rack over the cutout first.

> **WARNING:**
>
> If the power cords use overhead routing, then put them in a location where they will not fall or hit anyone.

4. For replacing PDU-B when there is no side panel access, and the rack does not have an InfiniBand cable harness:

> **Note:**
>
> Do not unstrap any cables attached to the cable arms.

   a. Unscrew the T25 screws holding the square cable arms to the rack.
   b. Move the InfiniBand cables to the middle, out of the way.
5. Unplug all power cables that connect the servers and switches to the PDU. Keep the power cables together in group bundles.
6. Remove the T30 screws from the top and bottom of the L-bracket, and note where the screws are used.
7. Note where the PDU sits in the rack frame.

   The PDU is typically an inch back from the rack frame, to allow access to the breaker switches.
8. Angle and maneuver the PDU out of the rack.
9. Hold the PDU or put it down, if there is enough room, while maneuvering the AC power cords through the rack. You might need to cut the cable ties that hold the AC cord flush with the bottom side of the PDU.
10. Pull the cords as near to the bottom or top of the rack as possible. There is more room between the servers to guide the outlet plug through the routing hole.
11. Remove the smaller Torx T25 screws, and loosen the nut on the top and bottom to remove the PDU from the L-bracket. You do not need to remove the nut.
12. Attach the L-bracket to the new PDU.
13. Put the new PDU next to the rack.
14. Route the AC cords through the rack to the outlets.

> **Note:**
>
> Do not cable-tie the AC cord to the new PDU.

15. Place the new PDU in the rack by angling and maneuvering it until the L-brackets rest on the top and bottom rails.
16. Line up the holes and slots so that the PDU is about an inch back from the rack frame.
17. Attach the power cords, using the labels on the cords as a guide. For example, G5-0 indicates PDU group 5 outlet 0 on the PDU.
18. Attach the InfiniBand cable holders, if you removed them in step 4. Oracle recommends that you first screw in the holders by hand to avoid stripping the screws.
19. Attach the AC power cords to the outlets.
20. Turn on the breakers.
21. Cable and program the PDU monitor for the network, as needed.

> **See Also:**
>
> *Oracle Sun Rack II Power Distribution Units User's Guide* for information about programming the PDU monitor at
>
> `http://docs.oracle.com/cd/E19844-01/html/E23956/index.html`

Replacing a Tape Drive

A tape drive can be replaced while Recovery Appliance is online.

> **Note:**
>
> The bridging tape drive is located in the top drive slot of the base module. The robot control is a SCSI Medium Changer device that appears as LUN 1 on a bridging tape drive. If the library is partitioned, there must be two tape drives in the base module, and each tape drive provides the robot control for its assigned partition. If the failed drive is the bridging drive, then the SL150 library must be taken offline to the host as replacing the tape drive will cause the host to lose connectivity to the SL150 library. If the tape drive is not a bridging drive, then it can be hot swapped.

To replace a tape drive:

1. Pause the Recovery Appliance `sbt_library`.

   ```
   exec dbms_ra.pause_sbt_library(lib_name=>'ROBOT0');
   ```

2. Quiesce activity for this tape drive.
3. Log into the SL150 remote interface using your browser.
4. (Optional) Enable the locate library indicator
5. Prepare the tape drive for removal:
   a. Click *Library* on the left of the remote interface.

b. Move the cursor to the drive you need to replace.

c. Click the drive icon and select *Remove Drive*.

d. Click OK in the confirmation dialog box. The physical LED at the rear of the drive tray lights to indicate the drive is ready for removal.

6. Remove the tape drive:

   a. Access the back of the library (open the rear door of the rack, if applicable).

   b. Locate the tape drive with the blue LED (indicates drive is ready for removal).

   c. Verify that the interface cables are labeled. Attach a label if necessary.

   d. Disconnect the cables from the jacks on the left side of the drive tray.

   e. Loosen the captive thumbscrews on the drive tray.

   f. Grasp the drive tray, pull it out of the library drive slot, and set it upright and flat on a static free surface.

> **Caution:**
>
> ESD damage. Do not touch any electronic components or contacts.

7. Remove the replacement drive from its packaging.

8. Replace the tape drive:

   a. Grasp the rear corners of the drive tray.

   b. Guide the front of the drive tray into the module drive slot.

   c. Push the drive tray completely into the drive slot.

   d. Verify that the LEDs are active on the rear of the drive tray.

   e. Tighten the captive thumbscrews firmly on each side of the drive tray to make sure there is no tray movement in any direction.

9. Push the Locate indicator on the robot CRU to extinguish the light, if applicable.

10. Confirm that the library recognizes and accounts for the drive from the Web GUI or front Op Panel.

11. Make sure the drive port is enabled. View the *Drive Properties* and change drive settings if appropriate.

12. Connect the interface and Ethernet cables to the proper jack on the left side of the drive tray.

13. Log out of the SL150 remote interface or return the touch screen to Home.

14. Run `obtool` to update serial number of the drive object.

```
obtool chdev -S <drive_object_replaced> | obtool chdev -S robot0_tape03
```

15. Resume ZDLRA `sbt_ library`.

```
exec dbms_ra.resume_sbt_library(lib_name=>'ROBOT0');
```

Many of the above steps come from How to Remove/Replace a Tape Drive CRU in the StorageTek SL150 Modular Tape Library [VCAP] (Doc ID 1473764.1).

Save Configuration before Updating Firmware

Preserve the existing Recovery Appliance configuration when updating Oracle ILOM system firmware.

You can update Oracle ILOM system firmware while the host is powered on. The Oracle ILOM firmware image includes firmware for the service processor (SP, Oracle ILOM) and the server's host components (FPGAs). The Oracle ILOM firmware update goes into effect immediately. However, the update of host components is deferred until the affected host is power cycled. Because Oracle ILOM can be updated while the host is powered on, this feature reduces the total system downtime.

Update the firmware using the Oracle ILOM command-line interface.

> **Note:**
>
> To save the configuration for the Recovery Appliance during a firmware upgrade, answer the prompts with y when asked `Preserve existing SP configuration (y/n)?`

1. Log in to Oracle ILOM with an account with Admin privileges.
2. Load the firmware image from the stored location using the `load -source` command followed by the directory path to the firmware image you want to install. Type:

   -> **`load -source`** **protocol**://server_IPaddress/<path_to_image>/<image.pkg>

   Where **protocol** can be: `http, https, ftp, tftp, sftp, scp`

   For example, if you are accessing the server through a tftp server with an IP address of 198.xxx.yyy.123 in a directory called `ilom/jdoe` and with the `<image.pkg>` named `firmware.pkg`, enter the following command:

   ```
   -> load -source tftp://198.xxx.yyy.123/tftpboot/ilom/jdoe/firmware.pkg
   ```

   The following information is displayed:

   ```
   An upgrade takes several minutes to complete. Oracle ILOM will enter a
   special mode to load new firmware. No other tasks can be performed in
   Oracle ILOM until the firmware upgrade is complete and Oracle ILOM is
   reset.

   You can choose to postpone the server BIOS upgrade until the next
   server power off. If you do not do that, you should perform a clean
   shutdown of the server before continuing.
   ```

3. Answer the following prompts:

   ```
   Are you sure you want to load the specified file? y

   Preserve existing SP configuration (y/n)? y
   ```

> **Note:**
>
> This prompt preserves your existing Oracle ILOM settings after the firmware update is complete. A y preserves the existing Recovery Appliance configuration.

```
Preserve existing BIOS configuration (y/n)? y
```

> **Note:**
>
> This prompt preserves your existing BIOS configuration settings after the firmware upgrade is complete.

```
Delay BIOS upgrade until the next poweroff or reset (y/n)? y
```

Answer "Y" (yes) to the Delay BIOS Upgrade question and, if the host is ON and there are host components to be updated, the host remains powered on and the host component updates are deferred until the next time the host powers off and powers on (next reset/reboot).

Answer "N" (no) to the Delay BIOS Upgrade question and, if the host is ON and there are host components to be updated, the host is forced OFF so that host component updates can be applied immediately. After Oracle ILOM reboots, the host is powered on automatically if it was forced off.

> **Note:**
>
> If the server has a pending BIOS upgrade, the power reset could take longer to complete. This is expected behavior, as it is necessary to power cycle the server to upgrade the BIOS firmware. If the upgrade includes an FPGA update, the process can take as long as 26 minutes to complete.

4. Wait for the Oracle ILOM status message to confirm that the process is complete.

> **Note:**
>
> For more details, see Update System Firmware Using Oracle ILOM

Resetting a Non-Responsive Oracle ILOM

The Oracle Integrated Lights Out Manager (Oracle ILOM) might become unresponsive. If this happens, then you must manually reset the Service Processor (SP) on Oracle ILOM.

The following procedures describe how to reset Oracle ILOM:

- Resetting Oracle ILOM Using SSH
- Resetting Oracle ILOM Using the Remote Console
- Resetting Oracle ILOM Using IPMItool

- Resetting Oracle ILOM By Removing Power

> **See Also:**
>
> Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation at
>
> http://docs.oracle.com/cd/E19860-01/E21549/bbgiedhj.html#z4000b491400243

Resetting Oracle ILOM Using SSH

To reset Oracle ILOM using SSH:

1. Connect to Oracle ILOM using SSH from another system.
2. Enter the following command at the ILOM prompt:

```
reset /SP
```

Resetting Oracle ILOM Using the Remote Console

If you cannot connect to Oracle ILOM using SSH, then log in to the remote console.

To reset Oracle ILOM using the remote console:

1. Log in to the Oracle ILOM remote console.
2. Select **Reset SP** from the Maintenance tab.
3. Click **Reset SP**.

Resetting Oracle ILOM Using IPMItool

If you cannot connect to Oracle ILOM using either SSH or the remote console, then use IPMItool.

To reset Oracle ILOM using IPMItool:

1. Log in to the local host or another host on the management network.
2. Use the following IPMItool command:
   - On the local host:

     ```
     $ ipmitool mc reset cold
     Sent cold reset command to MC
     ```

   - On another host:

     ```
     $ ipmitool -H ILOM_host_name -U ILOM_user mc reset cold
     Sent cold reset command to MC
     ```

     In the preceding command, *ILOM_host_name* is the host name being used, and *ILOM_user* is the user name for Oracle ILOM.

Resetting Oracle ILOM By Removing Power

If you cannot reset Oracle ILOM using the preceding options:

1. Unplug the server from the power supply.
2. Plug the server back into the power supply.

This action power cycles the server and Oracle ILOM.

Maintaining the Compute Servers

You do not need to shut down a compute server in Recovery Appliance to repair the physical disks. No downtime of the rack is required; however, individual servers might require downtime, and you might need to take them out of the cluster temporarily.

An LSI MegaRAID SAS 9261-8i disk controller manages the disk drives in each compute server. The disks have a RAID-5 configuration. Each compute server has four disk drives. One virtual drive comprises the RAID set.

> **See Also:**
>
> - "LED Status Descriptions" for information about the LEDs
> - "Parts for Compute Servers" for the repair procedures

Verifying the RAID Status of a Compute Server

Oracle recommends that you periodically verify the status of the compute server RAID devices. The impact is minimal. In contrast, the impact of corrective action varies depending on the specific issue uncovered, and can range from simple reconfiguration to an outage.

Log in to each compute server as `root` and perform the following procedure.

To verify the RAID status:

1. Check the current disk controller configuration:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -AdpAllInfo -aALL | grep "Device Present" -A 8

                Device Present
                ================
Virtual Drives    : 1
  Degraded        : 0
  Offline         : 0
Physical Devices  : 5
  Disks           : 4
  Critical Disks  : 0
  Failed Disks    : 0
```

   Verify that the output shows one virtual drive, none degraded or offline, five physical devices (one controller + four disks), four disks, and no critical or failed disks.

If the output is different, then investigate and correct the problem. Degraded virtual drives usually indicate absent or failed physical disks. Replace critical disks and failed disks immediately. Otherwise, you risk data loss if the number of working disks in the server is less than the number required to sustain normal operation.

2. Check the current virtual drive configuration:

```
# /opt/MegaRAID/MegaCli/MegaCli64 CfgDsply -aALL | grep "Virtual Drive:";
\
/opt/MegaRAID/MegaCli/MegaCli64 CfgDsply -aALL | grep "Number Of Drives";  \
/opt/MegaRAID/MegaCli/MegaCli64 CfgDsply -aALL | grep "^State"

Virtual Drive               : 0 (Target Id: 0)
Number Of Drives            : 4
State                       : Optimal
```

Verify that virtual device 0 has four drives, and the state is Optimal. If the output is different, then investigate and correct the problem.

3. Check the current physical drive configuration:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -PDList -aALL | grep "Firmware state"
Firmware state: Online, Spun Up
Firmware state: Online, Spun Up
Firmware state: Online, Spun Up
Firmware state: Online, Spun Up
```

Ensure that all drives are `Online, Spun Up`. If the output is different, then investigate and correct the problem.

If the output is different, then investigate and correct the problem. Degraded virtual drives usually indicate absent or failed physical disks. Replace critical disks and failed disks immediately. Otherwise, you risk data loss if the number of working disks in the server is less than the number required to sustain normal operation.

Reimaging a Compute Server

If a compute server is irretrievably damaged, then you must replace it and reimage the replacement server. During the reimaging procedure, the other compute servers in the cluster are available. When adding the new server to the cluster, you copy the software from a working compute server to the new server.

The following tasks describe how to reimage a compute server:

- "Contacting Oracle Support Services"
- "Downloading the Latest Release of the Cluster Verification Utility"
- "Removing the Failed Compute Server from the Cluster"
- "Preparing the USB Flash Drive for Imaging"
- "Copying the Image to the New Compute Server"
- "Configuring the Replacement Compute Server"
- "Preparing the Replacement Compute Server for the Cluster"
- "Applying Patch Bundles to a Replacement Compute Server"
- "Cloning the Oracle Grid Infrastructure"
- "Clone Oracle Database Homes to the Replacement Compute Server"

Contacting Oracle Support Services

Open a support request with Oracle Support Services. The support engineer identifies the failed server and sends you a replacement. The support engineer also asks for the output from the `imagehistory` command, run from a working compute server. The output provides a link to the `computeImageMaker` file that was used to image the original compute server, and is used to restore the system.

Downloading the Latest Release of the Cluster Verification Utility

The latest release of the cluster verification utility (`cluvfy`) is available from My Oracle Support Doc ID 316817.1.

Removing the Failed Compute Server from the Cluster

You must remove the failed compute server from Oracle Real Application Clusters (Oracle RAC).

In these steps, *working_server* is a working compute server in the cluster, *failed_server* is the compute server being replaced, and *replacement_server* is the new server.

To remove a failed compute server from the Oracle RAC cluster:

1. Log in to *working_server* as the `oracle` user.
2. Disable the listener that runs on the failed server:

   ```
   $ srvctl disable listener -n failed_server
   $ srvctl stop listener -n failed_server
   ```

3. Delete the Oracle home directory from the inventory:

   ```
   $ cd $ORACLE_HOME/oui/bin
   $ ./runInstaller -updateNodeList ORACLE_HOME= \
   /u01/app/oracle/product/12.1.0/dbhome_1 "CLUSTER_NODES=list_of_working_servers"
   ```

   In the preceding command, *list_of_working_servers* is a list of the compute servers that are still working in the cluster, such as `ra01db02`, `ra01db03`, and so on.

4. Verify that the failed server was deleted—that is, unpinned—from the cluster:

   ```
   $ olsnodes -s -t

   ra01db01     Inactive        Unpinned
   ra01db02        Active          Unpinned
   ```

5. Stop and delete the virtual IP (VIP) resources for the failed compute server:

   ```
   # srvctl stop vip -i failed_server-vip
   PRCC-1016 : failed_server-vip.example.com was already stopped

   # srvctl remove vip -i failed_server-vip
   Please confirm that you intend to remove the VIPs failed_server-vip (y/[n]) y
   ```

6. Delete the compute server from the cluster:

   ```
   # crsctl delete node -n failed_server
   CRS-4661: Node failed_server successfully deleted.
   ```

   If you receive an error message similar to the following, then relocate the voting disks.

```
CRS-4662: Error while trying to delete node ra01db01.
CRS-4000: Command Delete failed, or completed with errors.
```

To relocate the voting disks:

a. Determine the current location of the voting disks. The sample output shows that the current location is DBFS_DG.

```
# crsctl query css votedisk

##  STATE    File Universal Id          File Name                Disk
group
--  -----    -----------------          ---------
----------
1. ONLINE   123456789abab (o/192.168.73.102/DATA_CD_00_ra01cel07)
[DBFS_DG]
2. ONLINE   123456789cdcd (o/192.168.73.103/DATA_CD_00_ra01cel08)
[DBFS_DG]
3. ONLINE   123456789efef (o/192.168.73.100/DATA_CD_00_ra01cel05)
[DBFS_DG]
Located 3 voting disk(s).
```

b. Move the voting disks to another disk group:

```
# ./crsctl replace votedisk +DATA

Successful addition of voting disk 2345667aabbdd.
...
CRS-4266: Voting file(s) successfully replaced
```

c. Return the voting disks to the original location. This example returns them to DBFS_DG:

```
# ./crsctl replace votedisk +DBFS_DG
```

d. Repeat the `crsctl` command to delete the server from the cluster.

7. Update the Oracle inventory:

```
$ cd $ORACLE_HOME/oui/bin
$ ./runInstaller -updateNodeList ORACLE_HOME=/u01/app/12.1.0/grid \
  "CLUSTER_NODES=list_of_working_servers" CRS=TRUE
```

8. Verify that the server was deleted successfully:

```
$ cluvfy stage -post nodedel -n failed_server -verbose

Performing post-checks for node removal
Checking CRS integrity...
The Oracle clusterware is healthy on node "ra01db02"
CRS integrity check passed
Result:
Node removal check passed
Post-check for node removal was successful.
```

See Also:

Oracle Real Application Clusters Administration and Deployment Guide for information about deleting a compute server from a cluster

Preparing the USB Flash Drive for Imaging

Use a USB flash drive to copy the image to the new compute server.

To prepare the USB flash drive for use:

1. Insert a blank USB flash drive into a working compute server in the cluster.
2. Log in as the `root` user.
3. Unzip the `computeImage` file:

   ```
   # unzip computeImageMaker_release_LINUX.X64_release_date.platform.tar.zip

   # tar -xvf computeImageMaker_release_LINUX.X64_release_date.platform.tar
   ```

4. Load the image onto the USB flash drive:

   ```
   # cd dl360
   # ./makeImageMedia.sh -dualboot no
   ```

   The `makeImageMedia.sh` script prompts for information.

5. Remove the USB flash drive from the compute server.
6. Remove the unzipped `d1360` directory and the `computeImageMaker` file from the working compute server. The directory and file require about 2 GB of disk space.

Copying the Image to the New Compute Server

Before you perform the following procedure, replace the failed compute server with the new server. See Expanding a Recovery Appliance Rack with Additional Storage Servers.

To load the image onto the replacement server:

1. Insert the USB flash drive into the USB port on the replacement server.
2. Log in to the console through the service processor to monitor progress.
3. Power on the compute server either by physically pressing the power button or by using Oracle ILOM.
4. If you replaced the motherboard:
   a. Press F2 during BIOS
   b. Select **BIOS Setup**
   c. Set the USB flash drive first, and then the RAID controller.

   Otherwise, press F8 during BIOS, select the one-time boot selection menu, and choose the USB flash drive.
5. Allow the system to start.

   As the system starts, it detects the CELLUSBINSTALL media. The imaging process has two phases. Let both phases complete before proceeding to the next step.

   The first phase of the imaging process identifies any BIOS or firmware that is out of date, and upgrades the components to the expected level for the image. If any components are upgraded or downgraded, then the system automatically restarts.

   The second phase of the imaging process installs the factory image on the replacement compute server.

6. Remove the USB flash drive when the system prompts you.
7. Press Enter to power off the server.

Configuring the Replacement Compute Server

The replacement compute server does not have a host names, IP addresses, DNS, or NTP settings. This task describes how to configure the replacement compute server.

The information must be the same on all compute servers in Recovery Appliance. You can obtain the IP addresses from the DNS. You should also have a copy of the Installation Template from the initial installation.

To configure the replacement compute server:

1. Assemble the following information:
   - Name servers
   - Time zone, such as Americas/Chicago
   - NTP servers
   - IP address information for the management network
   - IP address information for the client access network
   - IP address information for the InfiniBand network
   - Canonical host name
   - Default gateway
2. Power on the replacement compute server. When the system starts, it automatically runs the configuration script and prompts for information.
3. Enter the information when prompted, and confirm the settings. The startup process then continues.

> **Note:**
>
> - If the compute server does not use all network interfaces, then the configuration process stops with a warning that some network interfaces are disconnected. It prompts whether to retry the discovery process. Respond with `yes` or `no`, as appropriate for the environment.
> - If bonding is used for the ingest network, then it is now set in the default active-passive mode.

Preparing the Replacement Compute Server for the Cluster

The initial installation of Recovery Appliance modified various files.

To modify the files on the replacement compute server:

1. Replicate the contents of the following files from a working compute server in the cluster:
   a. Copy the `/etc/security/limits.conf` file.

b. Merge the contents of the `/etc/hosts` files.

c. Copy the `/etc/oracle/cell/network-config/cellinit.ora` file.

d. Update the IP address with the IP address of the BONDIB0 interface on the replacement compute server.

e. Copy the `/etc/oracle/cell/network-config/cellip.ora` file.

f. Configure additional network requirements, such as 10 GbE.

g. Copy the `/etc/modprobe.conf` file.

h. Copy the `/etc/sysctl.conf` file.

i. Restart the compute server, so the network changes take effect.

2. Set up the Oracle software owner on the replacement compute server by adding the user name to one or more groups. The owner is usually the `oracle` user.

a. Obtain the current group information from a working compute server:

```
# id oracle
uid=1000(oracle) gid=1001(oinstall)
groups=1001(oinstall),1002(dba),1003(oper),1004(asmdba)
```

b. Use the `groupadd` command to add the group information to the replacement compute server. This example adds the groups identified in the previous step:

```
# groupadd –g 1001 oinstall
# groupadd –g 1002 dba
# groupadd –g 1003 oper
# groupadd –g 1004 asmdba
```

c. Obtain the current user information from a working compute server:

```
# id oracle uid=1000(oracle) gid=1001(oinstall) \
  groups=1001(oinstall),1002(dba),1003(oper),1004(asmdba)
```

d. Add the user information to the replacement compute server. This example adds the group IDs from the previous step to the `oracle` user ID:

```
# useradd -u 1000 -g 1001 -G 1001,1002,1003,1004 -m -d /home/oracle -s \
  /bin/bash oracle
```

e. Create the `ORACLE_BASE` and Grid Infrastructure directories. This example creates `/u01/app/oracle` and `/u01/app/12.1.0/grid`:

```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/12.1.0/grid
# chown -R oracle:oinstall /u01/app
```

f. Change the ownership of the `cellip.ora` and `cellinit.ora` files. The owner is typically `oracle:dba`.

```
# chown -R oracle:dba /etc/oracle/cell/network-config
```

g. Secure the restored compute server:

```
$ chmod u+x /opt/oracle.SupportTools/harden_passwords_reset_root_ssh
$ /opt/oracle.SupportTools/harden_passwords_reset_root_ssh
```

The compute server restarts.

h. Log in as the `root` user. When you are prompted for a new password, set it to match the `root` password of the other compute servers.

i. Set the password for the Oracle software owner. The owner is typically `oracle`.

```
# passwd oracle
```

3. Set up SSH for the `oracle` account:
   a. Change to the `oracle` account on the replacement compute server:

      ```
      # su - oracle
      ```

   b. Create the `dcli` group file on the replacement compute server, listing the servers in the Oracle cluster.
   c. Run the `setssh-Linux.sh` script on the replacement compute server. This example runs the script interactively:

      ```
      $ /opt/oracle.SupportTools/onecommand/setssh-Linux.sh -s
      ```

      The script prompts for the `oracle` password on the servers. The `-s` option causes the script to run in silent mode.
   d. Change to the `oracle` user on the replacement compute server:

      ```
      # su - oracle
      ```

   e. Verify SSH equivalency:

      ```
      $ dcli -g dbs_group -l oracle date
      ```

4. Set up or copy any custom login scripts from the working compute server to the replacement compute server:

   ```
   $ scp .bash* oracle@replacement_server:.
   ```

   In the preceding command, *replacement_server* is the name of the new server, such as `ra01db01`.

Applying Patch Bundles to a Replacement Compute Server

Oracle periodically releases software patch bundles for Recovery Appliance. If the working compute server has a patch bundle that is later than the release of the `computeImageMaker` file, then you must apply the patch bundle to the replacement compute server.

To determine if a patch bundle was applied, use the `imagehistory` command. Compare information on the replacement compute server to information on the working compute server. If the working database has a later release, then apply the storage server patch bundle to the replacement compute server.

Cloning the Oracle Grid Infrastructure

The following procedure describes how to clone the Oracle Grid infrastructure onto the replacement compute server. In the commands, *working_server* is a working compute server, and *replacement_server* is the replacement compute server.

To clone the Oracle Grid infrastructure:

1. Log in as `root` to a working compute server in the cluster.
2. Verify the hardware and operating system installation using the cluster verification utility (`cluvfy`):

   ```
   $ cluvfy stage -post hwos -n replacement_server,working_server –verbose
   ```

The phrase `Post-check for hardware and operating system setup was successful` should appear at the end of the report.

3. Verify peer compatibility:

```
$ cluvfy comp peer -refnode working_server -n replacement_server  \
  -orainv oinstall -osdba dba | grep -B 3 -A 2 mismatched
```

The following is an example of the output:

```
Compatibility check: Available memory [reference node: ra01db02]
Node Name Status Ref. node status Comment
------------ ---------------------- ---------------------- ----------
ra01db01 31.02GB (3.2527572E7KB) 29.26GB (3.0681252E7KB) mismatched
Available memory check failed
Compatibility check: Free disk space for "/tmp" [reference node: ra01db02]
Node Name Status Ref. node status Comment
------------ ---------------------- ---------------------- ----------
ra01db01 55.52GB (5.8217472E7KB) 51.82GB (5.4340608E7KB) mismatched
Free disk space check failed
```

If the only failed components are related to the physical memory, swap space, and disk space, then it is safe for you to continue.

4. Perform the requisite checks for adding the server:

   a. Ensure that the *GRID_HOME*`/network/admin/samples` directory has permissions set to 750.

   b. Validate the addition of the compute server:

```
$ cluvfy stage -ignorePrereq -pre nodeadd -n replacement_server \
-fixup -fixupdir  /home/oracle/fixup.d
```

   If the only failed component is related to swap space, then it is safe for you to continue.

   You might get an error about a voting disk similar to the following:

```
ERROR:
PRVF-5449 : Check of Voting Disk location "o/192.168.73.102/ \
DATA_CD_00_ra01cel07(o/192.168.73.102/DATA_CD_00_ra01cel07)" \
failed on the following nodes:
Check failed on nodes:
        ra01db01
        ra01db01:No such file or directory
...
PRVF-5431 : Oracle Cluster Voting Disk configuration check failed
```

   If this error occurs, then use the `-ignorePrereq` option when running the `addnode` script in the next step.

5. Add the replacement compute server to the cluster:

```
$ cd /u01/app/12.1.0/grid/addnode/
$ ./addnode.sh -silent "CLUSTER_NEW_NODES={replacement_server}" \
  "CLUSTER_NEW_VIRTUAL_HOSTNAMES={replacement_server-vip}"[-ignorePrereq]
```

The `addnode` script causes Oracle Universal Installer to copy the Oracle Clusterware software to the replacement compute server. A message like the following is displayed:

```
WARNING: A new inventory has been created on one or more nodes in this session.
However, it has not yet been registered as the central inventory of this
```

```
system. To register the new inventory please run the script at
'/u01/app/oraInventory/orainstRoot.sh' with root privileges on nodes
'ra01db01'. If you do not register the inventory, you may not be able to
update or patch the products you installed.

The following configuration scripts need to be executed as the "root" user in
each cluster node:

/u01/app/oraInventory/orainstRoot.sh #On nodes ra01db01

/u01/app/12.1.0/grid/root.sh #On nodes ra01db01
```

6. Run the configuration scripts:
   a. Open a terminal window.
   b. Log in as the `root` user.
   c. Run the scripts on each cluster server.

   After the scripts are run, the following message is displayed:

   ```
   The Cluster Node Addition of /u01/app/12.1.0/grid was successful.
   Please check '/tmp/silentInstall.log' for more details.
   ```

7. Run the `orainstRoot.sh` and `root.sh` scripts:

   ```
   # /u01/app/oraInventory/orainstRoot.sh
   Creating the Oracle inventory pointer file (/etc/oraInst.loc)
   Changing permissions of /u01/app/oraInventory.
   Adding read,write permissions for group.
   Removing read,write,execute permissions for world.
   Changing groupname of /u01/app/oraInventory to oinstall.
   The execution of the script is complete.

   # /u01/app/12.1.0/grid/root.sh
   ```

   Check the log files in `/u01/app/12.1.0/grid/install/` for the output of the `root.sh` script. The output file reports that the listener resource on the replaced compute server failed to start. This is an example of the expected output:

   ```
   /u01/app/12.1.0/grid/bin/srvctl start listener -n ra01db01 \
   ...Failed
   /u01/app/12.1.0/grid/perl/bin/perl \
   -I/u01/app/12.1.0/grid/perl/lib \
   -I/u01/app/12.1.0/grid/crs/install \
   /u01/app/12.1.0/grid/crs/install/rootcrs.pl execution failed
   ```

8. Reenable the listener resource that you stopped in "Removing the Failed Compute Server from the Cluster".

   ```
   # GRID_HOME/grid/bin/srvctl enable listener -l LISTENER \
     -n replacement_server

   # GRID_HOME/grid/bin/srvctl start listener -l LISTENER  \
     -n replacement_server
   ```

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide* for information about cloning

Clone Oracle Database Homes to the Replacement Compute Server

To clone the Oracle Database homes to the replacement server:

1. Add Oracle Database ORACLE_HOME to the replacement compute server:

   ```
   $ cd /u01/app/oracle/product/12.1.0/db_home/addnode/
   $ ./addnode.sh -silent "CLUSTER_NEW_NODES={replacement_server}" -ignorePrereq
   ```

   The addnode script causes Oracle Universal Installer to copy the Oracle Database software to the replacement compute server.

   ```
   WARNING: The following configuration scripts need to be executed as the "root"
   user in each cluster node.
   /u01/app/oracle/product/12.1.0/dbhome_1/root.sh #On nodes ra01db01
   To execute the configuration scripts:
   Open a terminal window.
   Log in as root.
   Run the scripts on each cluster node.
   ```

   After the scripts are finished, the following messages appear:

   ```
   The Cluster Node Addition of /u01/app/oracle/product/12.1.0/dbhome_1 was
   successful.
   Please check '/tmp/silentInstall.log' for more details.
   ```

2. Run the root.sh script on the replacement compute server:

   ```
   # /u01/app/oracle/product/12.1.0/dbhome_1/root.sh
   ```

   Check the /u01/app/oracle/product/12.1.0/dbhome_1/install/root_*replacement_server.company*.com_*date*.log file for the output of the script.

3. Ensure that the instance parameters are set for the replaced database instance. The following is an example for the CLUSTER_INTERCONNECTS parameter.

   ```
   SQL> SHOW PARAMETER cluster_interconnects

   NAME                                 TYPE        VALUE
   ------------------------------       --------    ------------------------
   cluster_interconnects                string

   SQL> ALTER SYSTEM SET cluster_interconnects='192.168.73.90' SCOPE=spfile
   SID='dbm1';
   ```

4. Validate the configuration files and correct them as necessary:
   - The ORACLE_HOME/dbs/init*SID*.ora file points to server parameter file (SPFILE) in the Oracle ASM shared storage.
   - The password file that is copied in the ORACLE_HOME/dbs directory has been changed to orapw*SID*.

5. Restart the database instance.

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide*

Maintaining the Storage Servers

This section describes how to perform maintenance on the storage servers. It contains the following topics:

- Shutting Down a Storage Server
- Enabling Network Connectivity Using the Diagnostics ISO

Shutting Down a Storage Server

When performing maintenance on a storage server, you might need to power down or restart the server. Before shutting down a storage server, verify that taking a server offline does not impact Oracle ASM disk group and database availability. Continued database availability depends on the level of Oracle ASM redundancy used on the affected disk groups, and the current status of disks in other storage servers that have mirror copies of the same data.

> **Caution:**
>
> - If a disk in a different cell fails while the cell undergoing maintenance is not completely back in service on the Recovery Appliance, a double disk failure can occur. If the Recovery Appliance is deployed with `NORMAL` redundancy for the `DELTA` disk group and if this disk failure is permanent, you will lose all backups on the Recovery Appliance.
> - Ensure that the cell undergoing maintenance is not offline for an extended period of time. Otherwise, a rebalance operation will occur and this will cause issues because of insufficient space for the operation to complete. By default, the rebalance operation begins 24 hours after the cell goes offline.

To power down a storage server:

1. Log in to the storage server as `root`.
2. (Optional) Keep the grid disks offline after restarting the storage server:

   ```
   CellCLI> ALTER GRIDDISK ALL INACTIVE
   ```

   Use this command when doing multiple restarts, or to control when the cell becomes active again. For example, so you can verify the planned maintenance activity was successful before the server is used.
3. Stop the cell services:

   ```
   CellCLI> ALTER CELL SHUTDOWN SERVICES ALL
   ```

The preceding command checks if any disks are offline, in predictive failure status, or must be copied to its mirror. If Oracle ASM redundancy is intact, then the command takes the grid disks offline in Oracle ASM, and stops the services.

The following error indicates that stopping the services might cause redundancy problems and force a disk group to dismount:

```
Stopping the RS, CELLSRV, and MS services...
The SHUTDOWN of ALL services was not successful.
CELL-01548: Unable to shut down CELLSRV because disk group DATA, RECO may be
forced to dismount due to reduced redundancy.
Getting the state of CELLSRV services... running
Getting the state of MS services... running
Getting the state of RS services... running
```

If this error occurs, then restore Oracle ASM disk group redundancy. Retry the command when the status is normal for all disks.

4. Shut down the server. See "Powering Down the Servers".
5. After you complete the maintenance procedure, power up the server. The services start automatically. During startup, all grid disks are automatically online in Oracle ASM.
6. Verify that all grid disks are online:

```
CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
```

Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

7. If you inactivated the grid disks in step 2, then reactivate them:

```
CellCLI> ALTER GRIDDISK ALL ACTIVE
```

If you skipped step 2, then the grid disks are activated automatically.

> **See Also:**
>
> My Oracle Support Doc ID 1188080.1, "Steps to shut down or reboot an Exadata storage cell without affecting ASM."

Enabling Network Connectivity Using the Diagnostics ISO

You might need to use the diagnostics ISO to access a storage server that fails to restart normally. After starting the server, you can copy files from the ISO to the server, replacing the corrupt files.

The ISO is located on all Recovery Appliance servers at `/opt/oracle.SupportTools/diagnostics.iso`.

> **Caution:**
>
> Use the diagnostics ISO only after other restart methods, such as using the USB drive, have failed. Contact Oracle Support for advise and guidance before starting this procedure.

To use the diagnostics ISO:

1. Enable a one-time CD-ROM boot in the service processor, using either the web interface or a serial console, such as Telnet or puTTY. For example, use this command from a serial console:

   ```
   set boot_device=cdrom
   ```

2. Mount a local copy of `diagnostics.iso` as a CD-ROM, using the service processor interface.
3. Use the `reboot` command to restart the server.
4. Log in to the server as the `root` user with the diagnostics ISO password.
5. To avoid pings:

   ```
   alias ping="ping -c"
   ```

6. Make a directory named `/etc/network`.
7. Make a directory named`/etc/network/if-pre-up.d`.
8. Add the following settings to the `/etc/network/interfaces` file, entering the actual IP address and netmask of the server, and the IP address of the gateway:

   ```
   iface eth0 inet static
   address IP address of server
   netmask netmask of server
   gateway gateway IP address of server
   ```

9. Start the eth0 interface:

   ```
   # ifup eth0
   ```

   Ignore any warning messages.
10. Use either FTP or the `wget` command to retrieve the files needed to repair the server.

Maintaining the Physical Disks of Storage Servers

This section contains the following topics:

- Reenabling a Rejected Physical Disk

> **See Also:**
>
> Oracle Maximum Availability Architecture (MAA) website at http://www.oracle.com/goto/maa for additional information about maintenance best practices

About System Disks and Data Disks

The first two disks of storage servers are system disks. Storage server software system software resides on a portion of each of the system disks. These portions on both system disks are referred to as the system area. The nonsystem area of the system disks, referred to as data partitions, is used for normal data storage. All other disks in a storage server are called data disks.

Monitoring the Status of Physical Disks

You can monitor a physical disk by checking its attributes with the CellCLI `LIST PHYSICALDISK` command. For example, a physical disk with a status of `failed` or `warning - predictive failure` is having problems and probably must be replaced. The disk firmware maintains the error counters, and marks a drive with `Predictive Failure` when internal thresholds are exceeded. The drive, not the server software, determines if it needs replacement.

The following list identifies the storage server physical disk statuses.

Physical Disk Status for Storage Servers

Physical Disk Status
normal
normal - dropped for replacement
normal - confinedOnline
normal - confinedOnline - dropped for replacement
not present
failed
failed - dropped for replacement
failed - rejected due to incorrect disk model
failed - rejected due to incorrect disk model - dropped for replacement
failed - rejected due to wrong slot
failed - rejected due to wrong slot - dropped for replacement
warning - confinedOnline
warning - confinedOnline - dropped for replacement
warning - peer failure
warning - poor performance
warning - poor performance - dropped for replacement
warning - poor performance, write-through caching
warning - predictive failure, poor performance
warning - predictive failure, poor performance - dropped for replacement

warning - predictive failure, write-through caching
warning - predictive failure
warning - predictive failure - dropped for replacement
warning - predictive failure, poor performance, write-through caching
warning - write-through caching

What Happens When Disk Errors Occur?

Oracle ASM performs bad extent repair for read errors caused by hardware errors. The disks stay online, and no alerts are sent.

When a disk fails:

- The Oracle ASM disks associated with it are dropped automatically with the `FORCE` option, and then an Oracle ASM rebalance restores data redundancy.
- The blue LED and the amber LED are turned on for the drive, indicating that disk replacement can proceed. The drive LED stays on solid. See "LED Status Descriptions" for information about LED status lights during predictive failure and poor performance.
- The server generates an alert, which includes specific instructions for replacing the disk. If you configured the system for alert notifications, then the alert is sent by email to the designated address.

When a disk has a faulty status:

- The Oracle ASM disks associated with the grid disks on the physical drive are dropped automatically.
- An Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.
- The blue LED is turned on for the drive, indicating that disk replacement can proceed.

When Oracle ASM gets a read error on a physically-addressed metadata block, it does not have mirroring for the blocks:

- Oracle ASM takes the disk offline.
- Oracle ASM drops the disk with the `FORCE` option.
- The storage server software sends an alert stating that the disk can be replaced.

See Also:

- "Replacing a Failed Physical Disk "
- "Replacing a Faulty Physical Disk"

About Detecting Underperforming Disks

ASR automatically identifies and removes a poorly performing disk from the active configuration. Recovery Appliance then runs a set of performance tests. When `CELLSRV` detects poor disk performance, the cell disk status changes to `normal -`

confinedOnline, and the physical disk status changes to warning - confinedOnline. Table 12-2 describes the conditions that trigger disk confinement:

Table 12-2 Alerts Indicating Poor Disk Performance

Alert Code	Cause
CD_PERF_HANG	Disk stopped responding
CD_PERF_SLOW_ABS	High service time threshold (slow disk)
CD_PERF_SLOW_RLTV	High relative service time threshold (slow disk)
CD_PERF_SLOW_LAT_WT	High latency on writes
CD_PERF_SLOW_LAT_RD	High latency on reads
CD_PERF_SLOW_LAT_RW	High latency on reads and writes
CD_PERF_SLOW_LAT_ERR	Frequent very high absolute latency on individual I/Os
CD_PERF_IOERR	I/O errors

If the problem is temporary and the disk passes the tests, then it is brought back into the configuration. If the disk does not pass the tests, then it is marked poor performance, and ASR submits a service request to replace the disk. If possible, Oracle ASM takes the grid disks offline for testing. Otherwise, the cell disk status stays at normal - confinedOnline until the disks can be taken offline safely. See "Removing an Underperforming Physical Disk".

The disk status change is recorded in the server alert history:

```
MESSAGE ID date_time info "Hard disk entered confinement status. The LUN
 n_m changed status to warning - confinedOnline. CellDisk changed status to normal
 - confinedOnline. Status: WARNING - CONFINEDONLINE  Manufacturer: name  Model
 Number: model  Size: size  Serial Number: serial_number  Firmware: fw_release
 Slot Number: m  Cell Disk: cell_disk_name  Grid Disk: grid disk 1, grid disk 2
      .
      .
      .
Reason for confinement: threshold for service time exceeded"
```

These messages are entered in the storage cell alert log:

```
CDHS: Mark cd health state change cell_disk_name  with newState HEALTH_BAD_
ONLINE pending HEALTH_BAD_ONLINE ongoing INVALID cur HEALTH_GOOD
Celldisk entering CONFINE ACTIVE state with cause CD_PERF_SLOW_ABS activeForced: 0
inactiveForced: 0 trigger HistoryFail: 0, forceTestOutcome: 0 testFail: 0
global conf related state: numHDsConf: 1 numFDsConf: 0 numHDsHung: 0 numFDsHung: 0
      .
      .
      .
```

About Rebalancing the Data

After you replace the physical disk, you must re-create the grid disks and cell disks that existed on the previous disk in that slot. If those grid disks were part of an Oracle ASM group, then add them back to the disk group, and rebalance the data, based on the disk group redundancy and the ASM_POWER_LIMIT parameter.

Oracle ASM rebalance occurs when dropping or adding a disk. To check the status of the rebalance:

- Did the rebalance operation run successfully?

  Check the Oracle ASM alert logs.

- Is the rebalance operation currently running?

  Check the `GV$ASM_OPERATION` view.

- Did the rebalance operation fail?

  Check the `V$ASM_OPERATION.ERROR` view.

You can perform rebalance operations from multiple disk groups on different Oracle ASM instances in the same cluster, if the failed physical disk contained ASM disks from multiple disk groups. One Oracle ASM instance can run one rebalance operation at a time. If all Oracle ASM instances are busy, then the rebalance operations are queued.

Monitoring Hard Disk Controller Write-Through Cache Mode

The hard disk controller on each storage server periodically performs a discharge and charge of the controller battery. During the operation, the write cache policy changes from write-back caching to write-through caching. Write-through cache mode is slower than write-back cache mode. However, write-back cache mode risks data loss if the storage server loses power or fails. The operation occurs every three months, for example, at 01:00 on the 17th day of January, April, July and October.

This example shows an informational alert that a storage server generates about the status of the caching mode for its logical drives:

```
HDD disk controller battery on disk contoller at adapter 0 is going into a learn
cycle. This is a normal maintenance activity that occurs quarterly and runs for
approximately 1 to 12 hours. The disk controller cache might go into WriteThrough
caching mode during the learn cycle. Disk write throughput might be temporarily
lower during this time. The message is informational only, no action is required.
```

Use the following commands to manage changes to the periodical write cache policy:

- To change the start time for the learn cycle, use a command like the following example:

  ```
  CellCLI> ALTER CELL bbuLearnCycleTime="2013-01-22T02:00:00-08:00"
  ```

  The time reverts to the default learn cycle time after the cycle completes.

- To see the time for the next learn cycle:

  ```
  CellCLI> LIST CELL ATTRIBUTES bbuLearnCycleTime
  ```

- To view the status of the battery:

  ```
  # /opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -GetBbuStatus -a0

  BBU status for Adapter: 0

  BatteryType: iBBU08
  Voltage: 3721 mV
  Current: 541 mA
  Temperature: 43 C

  BBU Firmware Status:
  Charging Status : Charging
  Voltage : OK
  ```

```
Temperature : OK
Learn Cycle Requested : No
Learn Cycle Active : No
Learn Cycle Status : OK
Learn Cycle Timeout : No
I2c Errors Detected : No
Battery Pack Missing : No
Battery Replacement required : No
Remaining Capacity Low : Yes
Periodic Learn Required : No
Transparent Learn : No

Battery state:

GasGuageStatus:
Fully Discharged : No
Fully Charged : No
Discharging : No
Initialized : No
Remaining Time Alarm : Yes
Remaining Capacity Alarm: No
Discharge Terminated : No
Over Temperature : No
Charging Terminated : No
Over Charged : No

Relative State of Charge: 7 %
Charger System State: 1
Charger System Ctrl: 0
Charging current: 541 mA
Absolute state of charge: 0 %
Max Error: 0 %

Exit Code: 0x00
```

Replacing a Failed Physical Disk

A physical disk outage can reduce performance and data redundancy. Therefore, you should replace a failed disk with a new disk as soon as possible.

To replace a disk when it fails:

1. Determine which disk failed.

```
CellCLI> LIST PHYSICALDISK WHERE diskType=HardDisk AND status=failed DETAIL

         name:                   28:5
         deviceId:               21
         diskType:               HardDisk
         enclosureDeviceId:      28
         errMediaCount:          0
         errOtherCount:          0
         foreignState:           false
         luns:                   0_5
         makeModel:              "SEAGATE ST360057SSUN600G"
         physicalFirmware:       0705
         physicalInterface:      sas
         physicalSerial:         A01BC2
         physicalSize:           558.9109999993816G
         slotNumber:             5
         status:                 failed
```

The slot number shows the location of the disk, and the status shows that the disk failed.

2. Ensure that the blue "OK to Remove" LED on the disk is lit, before you remove the disk.
3. Replace the physical disk on the storage server and wait three minutes. The physical disk is hot pluggable, and you can replace it with the power on.
4. Confirm that the disk is online and its status is `NORMAL`:

   ```
   CellCLI> LIST PHYSICALDISK WHERE name=28:5 ATTRIBUTES status
   ```

   When you replace a physical disk, the RAID controller must acknowledge the replacement disk before you can use it. Acknowledgment is quick.
5. Verify that the firmware is correct:

   ```
   ALTER CELL VALIDATE CONFIGURATION
   ```

   You can also check the `ms-odl.trc` file to confirm that the firmware was updated and the logical unit number (LUN) was rebuilt.
6. Re-create the grid disks and cell disks that existed on the previous disk in that slot. See "About Rebalancing the Data".

See Also:

- "Parts for Storage Servers"
- *Oracle Database Reference* about the `V$ASM_OPERATION` view

Replacing a Faulty Physical Disk

You might need to replace a physical disk because its status is `warning - predictive failure`. This status indicates that the physical disk will fail soon, and you should replace it at the earliest opportunity.

If the drive fails before you replace it, then see "Replacing a Failed Physical Disk ".

To replace a disk before it fails:

1. Identify the faulty disk:

```
CellCLI> LIST PHYSICALDISK WHERE diskType=HardDisk AND status= \
       "warning - predictive failure" DETAIL

         name:                   28:3
         deviceId:               19
         diskType:               HardDisk
         enclosureDeviceId:      28
         errMediaCount:          0
         errOtherCount:          0
         foreignState:           false
         luns:                   0_3
         makeModel:              "SEAGATE ST360057SSUN600G"
         physicalFirmware:       0705
         physicalInterface:      sas
```

```
physicalSerial:         E07L8E
physicalSize:           558.9109999993816G
slotNumber:             3
status:                 warning - predictive failure
```

In the sample output from the previous command, the slot number shows the location of the disk, and the status shows that the disk is expected to fail.

2. Ensure that the blue "OK to Remove" LED on the disk is lit, before you remove the disk.
3. Wait while the affected Oracle ASM disks are dropped. To check the status, query the `V$ASM_DISK_STAT` view on the Oracle ASM instance.

> **Caution:**
>
> The disks in the first two slots are system disks, which store the operating system and the Recovery Appliance storage server software. One system disk must be in working condition for the server to operate.
>
> Before replacing the other system disk, wait until `ALTER CELL VALIDATE CONFIGURATION` shows no RAID `mdadm` errors. This output indicates that the system disk resynchronization is complete.

> **See Also:**
>
> *Oracle Database Reference* for information about querying the `V$ASM_DISK_STAT` view

4. Replace the physical disk on the storage server and wait three minutes. The physical disk is hot pluggable, and you can replace it when the power is on.
5. Confirm that the disk is online and its status is `NORMAL`:

   ```
   CellCLI> LIST PHYSICALDISK WHERE name=28:5 ATTRIBUTES status
   ```

   When you replace a physical disk, the RAID controller must acknowledge the replacement disk before you can use it. Acknowledgment is quick.
6. Verify that the firmware is correct:

   ```
   ALTER CELL VALIDATE CONFIGURATION
   ```

7. Re-create the grid disks and cell disks that existed on the previous disk in that slot. See "About Rebalancing the Data".

> **See Also:**
>
> - "Parts for Storage Servers"
> - *Oracle Database Reference* for information about the `V$ASM_OPERATION` view

Removing an Underperforming Physical Disk

A bad physical disk can degrade the performance of other good disks. You should remove the bad disk from the system.

To remove a physical disk after identifying the bad disk:

1. Illuminate the physical drive service LED to identify the drive to be replaced:

   ```
   cellcli -e 'alter physicaldisk disk_name serviceled on'
   ```

   In the preceding command, *disk_name* is the name of the physical disk to be replaced, such as `20:2`.

2. Identify all grid disks on the bad disk, and direct Oracle ASM to stop using them:

   ```
   ALTER DISKGROUP diskgroup_name DROP DISK asm_disk_name
   ```

3. Ensure that the blue "OK to Remove" LED on the disk is lit.
4. Query the `V$ASM_DISK_STAT` view to ensure that the Oracle ASM disks affected by the bad disk were dropped successfully.
5. Remove the bad disk.

   An alert is sent when the disk is removed.

6. When a new disk is available, install it in the system. The cell disks and grid disks are created automatically on the new physical disk.
7. Confirm that the disk is online and its status is `NORMAL`:

   ```
   CellCLI> LIST PHYSICALDISK WHERE name=28:5 ATTRIBUTES status
   ```

   When you replace a physical disk, the RAID controller must acknowledge the replacement disk before you can use it. Acknowledgment is quick.

> **See Also:**
>
> "About Detecting Underperforming Disks"

Moving All Drives from One Storage Server to Another

You might need to move all drives from one storage server to another storage server. This situation might occur when a chassis-level component fails, such as a motherboard or Oracle ILOM, or when you are troubleshooting a hardware problem.

To move the drives between storage servers:

1. Back up the files in the following directories:
   - `/etc/hosts`
   - `/etc/modprobe.conf`
   - `/etc/sysconfig/network`
   - `/etc/sysconfig/network-scripts`

2. Inactivate all grid disks and shut down the storage server. See "Shutting Down a Storage Server".
3. Ensure that the Oracle ASM `disk_repair_time` attribute is set long enough, so that Oracle ASM does not drop the disks before you can activate the grid disks in another storage server.
4. Move the physical disks, flash disks, disk controller, and USB flash drive from the original storage server to the new storage server.

   Caution:

   - Ensure that the first two disks, which are the system disks, are in the same, first two slots. Otherwise, the storage server will function improperly.
   - Ensure that the flash cards are installed in the same PCIe slots as in the original storage server.

5. Power on the new storage server. You can either use the service processor interface or press the power button.
6. Log in to the console using the service processor.
7. Check the files in the following directories. Restore corrupt files from the backups.
   - `/etc/hosts`
   - `/etc/modprobe.conf`
   - `/etc/sysconfig/network`
   - `/etc/sysconfig/network-scripts`
8. Use the `ifconfig` command to retrieve the new MAC addresses for eth0, eth1, eth2, and eth3. This example shows that the eth0 MAC address (HWaddr) is 00:14:4F:CA:D9:AE.

```
# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:14:4F:CA:D9:AE
          inet addr:10.204.74.184  Bcast:10.204.75.255  Mask:255.255.252.0
          inet6 addr: fe80::214:4fff:feca:d9ae/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:141455 errors:0 dropped:0 overruns:0 frame:0
          TX packets:6340 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:9578692 (9.1 MiB)  TX bytes:1042156 (1017.7 KiB)
          Memory:f8c60000-f8c80000
```

9. In the `/etc/sysconfig/network-scripts` directory, edit the following files to change HWADDR to the value returned in step 8:

   ```
   ifcfg-eth0
   ifcfg-eth1
   ifcfg-eth2
   ifcfg-eth3
   ```

   The following example shows the edited `ifcfg-eth0` file:

```
#### DO NOT REMOVE THESE LINES ####
#### %GENERATED BY CELL% ####
```

```
DEVICE=eth0
BOOTPROTO=static
ONBOOT=yes
IPADDR=10.204.74.184
NETMASK=255.255.252.0
NETWORK=10.204.72.0
BROADCAST=10.204.75.255
GATEWAY=10.204.72.1
HOTPLUG=no
IPV6INIT=no
HWADDR=00:14:4F:CA:D9:AE
```

10. Restart the storage server.
11. Activate the grid disks:

    ```
    CellCLI> ALTER GRIDDISK ALL ACTIVE
    ```

    If the Oracle ASM disks were not dropped, then they go online automatically and start being used.
12. Validate the configuration:

    ```
    CellCLI> ALTER CELL VALIDATE CONFIGURATION
    ```
13. Activate Oracle ILOM for ASR.

Removing and Replacing the Same Physical Disk

If you remove the wrong physical disk and replace it, then Recovery Appliance automatically adds the disk back in the Oracle ASM disk group, and resynchronizes its data.

> **Note:**
>
> When replacing a faulty or failed disk, look for a lit LED on the disk. The LED is lit to help you locate the bad disk.

Reenabling a Rejected Physical Disk

Recovery Appliance rejects a physical disk when it is in the wrong slot.

> **Caution:**
>
> Reenabling a physical disk removes all data stored on it.

- To reenable a rejected physical disk, replace *hard_disk_name* and *hard_disk_id* with the appropriate values in this command:

  ```
  CellCLI> ALTER PHYSICALDISK hard_disk_name/hard_disk_id reenable force
  Physical disk hard_disk_name/hard_disk_id  was reenabled.
  ```

Maintaining the Flash Disks of Storage Servers

This section describes how to perform maintenance on flash disks. It contains the following topics:

About the Flash Disks

Recovery Appliance mirrors data across storage servers, and sends write operations to at least two storage servers. If a flash card in one storage server has problems, then Recovery Appliance services the read and write operations using the mirrored data in another storage server. Service is not interrupted.

If a flash card fails, then the storage server software identifies the data in the flash cache by reading the data from the surviving mirror. It then writes the data to the server with the failed flash card. When the failure occurs, the software saves the location of the data lost in the failed flash cache. Resilvering then replaces the lost data with the mirrored copy. During resilvering, the grid disk status is `ACTIVE -- RESILVERING WORKING`.

Each storage server has four PCIe cards. Each card has four flash disks (FDOMs) for a total of 16 flash disks. The four PCIe cards are located in PCI slot numbers 1, 2, 4, and 5.

To identify a failed flash disk, use the following command:

```
CellCLI> LIST PHYSICALDISK WHERE DISKTYPE=flashdisk AND STATUS=failed DETAIL

         name:                   FLASH_5_3
         diskType:               FlashDisk
         luns:                   5_3
         makeModel:              "Sun Flash Accelerator F40 PCIe Card"
         physicalFirmware:       TI35
         physicalInsertTime:     2012-07-13T15:40:59-07:00
         physicalSerial:         5L002X4P
         physicalSize:           93.13225793838501G
         slotNumber:             "PCI Slot: 5; FDOM: 3"
         status:                 failed
```

The card `name` and `slotNumber` attributes show the PCI slot and the FDOM number.

When the server software detects a failure, it generates an alert that indicates that the flash disk, and the LUN on it, failed. The alert message includes the PCI slot number of the flash card and the exact FDOM number. These numbers uniquely identify the field replaceable unit

(FRU). If you configured the system for alert notification, then the alert is sent to the designated address in an email message.

A flash disk outage can reduce performance and data redundancy. Replace the failed disk at the earliest opportunity. If the flash disk is used for flash cache, then the effective cache size for the server is reduced. If the flash disk is used for flash log, then the flash log is disabled on the disk, thus reducing the effective flash log size. If the flash disk is used for grid disks, then the Oracle ASM disks associated with them are automatically dropped with the `FORCE` option from the Oracle ASM disk group, and an Oracle ASM rebalance starts to restore the data redundancy.

See Also:

- "Parts for Storage Servers" for part number information and a link to the service guide
- *Oracle Database Reference* for information about the `V$ASM_OPERATION` view
- *Sun Flash Accelerator F80 PCIe Card User's Guide* at

  `http://docs.oracle.com/cd/E41278_01/html/E41251/index.html`

Faulty Status Indicators

The following status indicators generate an alert. The alert includes specific instructions for replacing the flash disk. If you configured the system for alert notifications, then the alerts are sent by email message to the designated address.

warning - peer failure
One of the flash disks on the same Sun Flash Accelerator PCIe card failed or has a problem. For example, if FLASH5_3 fails, then FLASH5_0, FLASH5_1, and FLASH5_2 have peer failure status:

```
CellCLI> LIST PHYSICALDISK
         36:0            L45F3A          normal
         36:1            L45WAE          normal
         36:2            L45WQW          normal
          .
          .
          .
         FLASH_5_0       5L0034XM        warning - peer failure
         FLASH_5_1       5L0034JE        warning - peer failure
         FLASH_5_2       5L002WJH        warning - peer failure
         FLASH_5_3       5L002X4P        failed
```

warning - predictive failure
The flash disk will fail soon, and should be replaced at the earliest opportunity. If the flash disk is used for flash cache, then it continues to be used as flash cache. If the flash disk is used for grid disks, then the Oracle ASM disks associated with these grid disks are automatically dropped, and Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.

When one flash disk has predictive failure status, then the data is copied. If the flash disk is used for write back flash cache, then the data is flushed from the flash disks to the grid disks.

warning - poor performance
The flash disk demonstrates extremely poor performance, and should be replaced at the earliest opportunity. If the flash disk is used for flash cache, then flash cache is dropped from this disk, thus reducing the effective flash cache size for the storage server. If the flash disk is used for grid disks, then the Oracle ASM disks associated with the grid disks on this flash disk are automatically dropped with the `FORCE` option, if possible. If `DROP...FORCE` cannot succeed because of offline partners, then the grid disks are dropped normally, and Oracle ASM rebalance relocates the data from the poor performance disk to the other disks.

warning - write-through caching
The capacitors used to support data cache on the PCIe card failed, and the card should be replaced as soon as possible.

Identifying Flash Disks in Poor Health

To identify a flash disk with a particular health status, use the `LIST PHYSICALDISK` command. This example queries for the `warning - predictive failure` status:

```
CellCLI> LIST PHYSICALDISK WHERE DISKTYPE=flashdisk AND STATUS=  \
'warning - predictive failure' DETAIL


         name:                   FLASH_5_3
         diskType:               FlashDisk
         luns:                   5_3
         makeModel:              "Sun Flash Accelerator F40 PCIe Card"
         physicalFirmware:       TI35
         physicalInsertTime:     2012-07-13T15:40:59-07:00
         physicalSerial:         5L002X4P
         physicalSize:           93.132257938385O1G
         slotNumber:             "PCI Slot: 1; FDOM: 2"
         status:                 warning - predictive failure
```

Identifying Underperforming Flash Disks

ASR automatically identifies and removes a poorly performing disk from the active configuration. Recovery Appliance then runs a set of performance tests. When `CELLSRV` detects poor disk performance, the cell disk status changes to `normal - confinedOnline`, and the physical disk status changes to `warning - confinedOnline`. Table 12-2 describes the conditions that trigger disk confinement. The conditions are the same for both physical and flash disks.

If the problem is temporary and the disk passes the tests, then it is brought back into the configuration. If the disk does not pass the tests, then it is marked `poor performance`, and ASR submits a service request to replace the disk. If possible, Oracle ASM takes the grid disks offline for testing. Otherwise, the cell disk status stays at `normal - confinedOnline` until the disks can be taken offline safely.

The disk status change is recorded in the server alert history:

```
MESSAGE ID date_time info "Hard disk entered confinement status. The LUN
 n_m changed status to warning - confinedOnline. CellDisk changed status to normal
 - confinedOnline. Status: WARNING - CONFINEDONLINE  Manufacturer: name  Model
 Number: model  Size: size  Serial Number: serial_number  Firmware: fw_release
```

```
 Slot Number: m  Cell Disk: cell_disk_name  Grid Disk: grid disk 1, grid disk 2
 ... Reason for confinement: threshold for service time exceeded"
```

These messages are entered in the storage cell alert log:

```
CDHS: Mark cd health state change cell_disk_name  with newState HEALTH_BAD_
ONLINE pending HEALTH_BAD_ONLINE ongoing INVALID cur HEALTH_GOOD
Celldisk entering CONFINE ACTIVE state with cause CD_PERF_SLOW_ABS activeForced:
0
inactiveForced: 0 trigger HistoryFail: 0, forceTestOutcome: 0 testFail: 0
global conf related state: numHDsConf: 1 numFDsConf: 0 numHDsHung: 0 numFDsHung:
0
     .
     .
     .
```

When Is It Safe to Replace a Faulty Flash Disk?

When the server software detects a predictive or peer failure in a flash disk used for write back flash cache, and only one FDOM is bad, then the server software resilvers the data on the bad FDOM, and flushes the data on the other three FDOMs. If there are valid grid disks, then the server software initiates an Oracle ASM rebalance of the disks. You cannot replace the bad disk until the tasks are completed and an alert indicates that the disk is ready.

An alert is sent when the Oracle ASM disks are dropped, and you can safely replace the flash disk. If the flash disk is used for write-back flash cache, then wait until none of the grid disks are cached by the flash disk.

Replacing a Failed Flash Disk

> **Caution:**
>
> The PCIe cards are not hot pluggable; you must power down a storage server before replacing the flash disks or cards.

Before you perform the following procedure, shut down the server. See "Shutting Down a Storage Server".

To replace a failed flash disk:

1. Replace the failed flash disk. Use the PCI number and FDOM number to locate the failed disk. A white cell LED is lit to help you locate the affected server.
2. Power up the server. The services start automatically. As part of the server startup, all grid disks are automatically online in Oracle ASM.
3. Verify that all grid disks are online:

   ```
   CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
   ```

   Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

See Also:

- "Parts for Storage Servers" for part numbers and links to the service guide
- *Oracle Database Reference* for information about the `V$ASM_OPERATION` view
- *Sun Flash Accelerator F80 PCIe Card User's Guide* at `http://docs.oracle.com/cd/E41278_01/html/E41251/index.html`

Replacing a Faulty Flash Disk

Caution:

The PCIe cards are not hot pluggable; you must power down a storage server before replacing the flash disks or cards.

Before you perform the following procedure, review the "When Is It Safe to Replace a Faulty Flash Disk?" topic.

To replace a faulty flash disk:

1. Use the following command to check the `cachedBy` attribute of all grid disks.

   ```
   CellCLI> LIST GRIDDISK ATTRIBUTES name, cachedBy
   ```

   The cell disk on the flash disk should not appear in any grid disk `cachedBy` attribute. If the flash disk is used for both grid disks and flash cache, then wait until receiving the alert, and the cell disk is not shown in any grid disk `cachedBy` attribute.

2. Stop all services:

   ```
   CellCLI> ALTER CELL SHUTDOWN SERVICES ALL
   ```

   The preceding command checks if any disks are offline, in predictive failure status, or must be copied to a mirror. If Oracle ASM redundancy is intact, then the command takes the grid disks offline in Oracle ASM, and then stops the services.

   The following error indicates that it might be unsafe to stop the services, because stopping them might force a disk group to dismount:

   ```
   Stopping the RS, CELLSRV, and MS services...
   The SHUTDOWN of ALL services was not successful.
   CELL-01548: Unable to shut down CELLSRV because disk group DATA, RECO may be
   forced to dismount due to reduced redundancy.
   Getting the state of CELLSRV services... running
   Getting the state of MS services... running
   Getting the state of RS services... running
   ```

   If this error occurs, then restore Oracle ASM disk group redundancy, and retry the command when the disk status is normal for all disks.

3. Shut down the server.

   See "Shutting Down a Storage Server".

4. Replace the failed flash disk. Use the PCI number and FDOM number to locate the failed disk. A white cell LED is lit to help you locate the affected server.
5. Power up the server. The services start automatically. As part of the server startup, all grid disks are automatically online in Oracle ASM.
6. Verify that all grid disks are online:

   ```
   CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
   ```

   Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

The system automatically uses the new flash disk, as follows:

- If the flash disk is used for flash cache, then the effective cache size increases.
- If the flash disk is used for grid disks, then the grid disks are re-created on the new flash disk.
- If the grid disks were part of an Oracle ASM disk group, then they are added back to the disk group. The data is rebalanced on them, based on the disk group redundancy and the `ASM_POWER_LIMIT` parameter.

Removing an Underperforming Flash Disk

A bad flash disk can degrade the performance of other good flash disks. You should remove a bad flash disk. See "Identifying Underperforming Flash Disks".

To remove an underperforming flash drive:

1. If the flash disk is used for flash cache:
   a. Ensure that data not synchronized with the disk (dirty data) is flushed from flash cache to the grid disks:

      ```
      CellCLI> ALTER FLASHCACHE ... FLUSH
      ```

   b. Disable the flash cache and create a new one. Do not include the bad flash disk when creating the flash cache.

      ```
      CellCLI > DROP FLASHCACHE
      CellCLI > CREATE FLASHCACHE CELLDISK='fd1,fd2,fd3,fd4, ...'
      ```

2. If the flash disk is used for grid disks, then direct Oracle ASM to stop using the bad disk immediately:

   ```
   SQL> ALTER DISKGROUP diskgroup_name DROP DISK asm_disk_name FORCE
   ```

   Offline partners might cause the `DROP` command with the `FORCE` option to fail. If the previous command fails, do one of the following:

   - Restore Oracle ASM data redundancy by correcting the other server or disk failures. Then retry the `DROP...FORCE` command.
   - Direct Oracle ASM to rebalance the data off the bad disk:

     ```
     SQL> ALTER DISKGROUP diskgroup_name DROP DISK asm_disk_name  NOFORCE
     ```

3. Wait until the Oracle ASM disks associated with the bad flash disk are dropped successfully. The storage server software automatically sends an alert when it is safe to replace the flash disk.
4. Stop the services:

   ```
   CellCLI> ALTER CELL SHUTDOWN SERVICES ALL
   ```

The preceding command checks if any disks are offline, in predictive failure status, or must be copied to its mirror. If Oracle ASM redundancy is intact, then the command takes the grid disks offline in Oracle ASM, and stops the services.

The following error indicates that stopping the services might cause redundancy problems and force a disk group to dismount:

```
Stopping the RS, CELLSRV, and MS services...
The SHUTDOWN of ALL services was not successful.
CELL-01548: Unable to shut down CELLSRV because disk group DATA, RECO may be
forced to dismount due to reduced redundancy.
Getting the state of CELLSRV services... running
Getting the state of MS services... running
Getting the state of RS services... running
```

If this error occurs, then restore Oracle ASM disk group redundancy. Retry the command when the status is normal for all disks.

5. Shut down the server. See "Shutting Down a Storage Server".
6. Remove the bad flash disk, and replace it with a new flash disk.
7. Power up the server. The services are started automatically. As part of the server startup, all grid disks are automatically online in Oracle ASM.
8. Add the new flash disk to flash cache:

```
CellCLI> DROP FLASHCACHE
CellCLI> CREATE FLASHCACHE ALL
```

9. Verify that all grid disks are online:

```
CellCLI> LIST GRIDDISK ATTRIBUTES asmmodestatus
```

Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

The flash disks are added as follows:

- If the flash disk is used for grid disks, then the grid disks are re-created on the new flash disk.
- If these grid disks were part of an Oracle ASM disk group and `DROP...FORCE` was used in Step 2, then they are added back to the disk group and the data is rebalanced on based on disk group redundancy and the `ASM_POWER_LIMIT` parameter.
- If `DROP...NOFORCE` was used in Step 2, then you must manually add the grid disks back to the Oracle ASM disk group.

About Write-Back Flash Cache

You cannot modify the write-back flash cache settings on Recovery Appliance.

Replacing a Disk Controller Battery Backup Unit

The disk controller battery backup unit (disk controller BBU) resides on a drive tray in the compute and storage servers. You can replace the disk controller BBU without downtime. The following procedures describe how to replace the disk controller BBU:

- Replacing a Disk Controller BBU on a Compute Server
- Replacing a Disk Controller BBU on a Storage Server

> **Note:**
>
> The procedures in this section do not apply to on-controller battery backup units. Replacement of those units require a system shutdown, because the system must be opened to access the controller card.

Replacing a Disk Controller BBU on a Compute Server

The following procedure describes how to replace a disk controller BBU on a compute server:

1. Drop the disk controller BBU for replacement:

   ```
   # /opt/oracle.cellos/compmon/exadata_mon_hw_asr.pl -drop_bbu_for_replacement
   ```

2. Verify that the disk controller BBU has been dropped for replacement:

   ```
   # /opt/oracle.cellos/compmon/exadata_mon_hw_asr.pl -list_bbu_status

   BBU status: dropped for replacement.
   ```

3. Replace the disk controller BBU by releasing the drive caddy and slowly pulling out the tray, and then sliding the replacement tray into the slot. The disk controller BBU is located in slot 7.

4. Verify that the new disk controller BBU has been detected. It may take several minutes.

   ```
   # /opt/oracle.cellos/compmon/exadata_mon_hw_asr.pl -list_bbu_status

   BBU status: present
   ```

5. Verify that the current logical disk drive cache policy uses `writeback` mode.

   ```
   # /opt/MegaRAID/MegaCli/MegaCli64 -ldinfo -lall -a0 | egrep         \
   'Default Cache|Current Cache'
   Default Cache Policy: WriteBack, ReadAheadNone, Direct, No Write Cache if
   Bad BBU
   Current Cache Policy: WriteBack, ReadAheadNone, Direct, No Write Cache if
   Bad BBU
   ```

   If the cache policy is not `writeback`, then go to step 6. Otherwise, go to step 7.

6. Verify that the battery state is `Operational`. This step is required only when the cache policy output from step 5 is not `writeback`.

   ```
   # /opt/MegaRAID/MegaCli/MegaCli64 -adpbbucmd -getbbustatus -a0|grep Battery
   BatteryType: iBBU08
   Battery State : Operational
   Battery Pack Missing : No
   Battery Replacement required : No
   ```

   If the battery state is not `Operational`, then investigate and correct the problem.

7. Perform battery checks as described in My Oracle Support Doc ID 1274318.1. If the checks return unexpected results, then refer to the note for additional information and instructions.

8. (Optional) Use the `exachk` tool to verify the health of the system. See My Oracle Support Doc ID 1070954.1.

Replacing a Disk Controller BBU on a Storage Server

To replace a disk controller BBU on a storage server:

1. Drop the disk controller BBU for replacement using the following command:

   ```
   # cellcli -e alter cell bbu drop for replacement
   ```

2. Verify that the disk controller BBU has been dropped for replacement using the following command:

   ```
   # cellcli -e list cell attributes bbustatus

   BBU status: dropped for replacement.
   ```

3. Replace the disk controller BBU by releasing the drive caddy and slowly pulling out the tray, and then sliding the replacement tray into the slot. The disk controller BBU is located in rear slot 1 of the server.

4. Verify that the disk controller BBU battery state is operational.

   ```
   # cellcli -e list cell attributes bbustatus

   BBU status: normal
   ```

5. Perform battery checks as described in My Oracle Support Doc ID 1274318.1. If the checks return unexpected results, then refer to the note for additional information and instructions.

6. (Optional) Use the `exachk` tool to verify the health of the system. The tool is available in My Oracle Support Doc ID 1070954.1.

Using the Storage Server Rescue Procedure

Each storage server maintains a copy of the software on the USB stick. Whenever the system configuration changes, the server updates the USB stick. You can use this USB stick to recover the server after a hardware replacement or a software failure. You restore the system when the system disks fail, the operating system has a corrupt file system, or the boot area is damaged. You can replace the disks, cards, CPU, memory, and so forth, and recover the server. You can insert the USB stick in a different server, and it will duplicate the old server.

If only one system disk fails, then use CellCLI commands to recover. In the rare event that both system disks fail simultaneously, then use the rescue functionality provided on the storage server CELLBOOT USB flash drive.

This section contains the following topics:

First Steps Before Rescuing the Storage Server

Before rescuing a storage server, you must take steps to protect the data that is stored on it. Those steps depend on whether the system is set up with normal redundancy or high redundancy.

If the Server Has Normal Redundancy

If you are using normal redundancy, then the server has one mirror copy. The data could be irrecoverably lost, if that single mirror also fails during the rescue procedure.

Oracle recommends that you duplicate the mirror copy:

1. Make a complete backup of the data in the mirror copy.
2. Take the mirror copy server offline immediately, to prevent any new data changes to it before attempting a rescue.

This procedure ensures that all data residing on the grid disks on the failed server and its mirror copy is inaccessible during the rescue procedure.

The Oracle ASM disk repair timer has a default repair time of 3.6 hours. If you know that you cannot perform the rescue procedure within that time frame, then use the Oracle ASM rebalance procedure to rebalance the disks until you can do the rescue procedure.

> **See Also:**
>
> *Oracle Exadata Storage Server Software User's Guide* for information about resetting the timer

If the Server Has High Redundancy

When the server has high redundancy disk groups, so that Oracle ASM has multiple mirror copies for all the grid disks of the failed server, then take the failed cell offline. After Oracle ASM times out, it automatically drops the grid disks on the failed server, and starts rebalancing the data using mirror copies.

The default time out is two hours. If the server rescue takes more than two hours, then you must re-create the grid disks on the rescued cells in Oracle ASM.

About the Rescue Procedure

Note the following before using the rescue procedure:

- The rescue procedure can rewrite some or all of the disks in the cell. If this happens, then you might lose all the content of those disks without the possibility of recovery. Ensure that you complete the appropriate preliminary steps before starting the rescue. See "If the Server Has Normal Redundancy" or "If the Server Has High Redundancy".

- Use extreme caution when using this procedure, and pay attention to the prompts. Ideally, use the rescue procedure only with assistance from Oracle Support Services, and when you can afford to lose the data on some or all of the disks.
- The rescue procedure does not destroy the contents of the data disks or the contents of the data partitions on the system disks, unless you explicitly choose to do so during the rescue procedure.
- The rescue procedure restores the storage server software to the same release, including any patches that existed on the server during the last successful boot.
- The rescue procedure does *not* restore these configuration settings:
  - Server configurations, such as alert configurations, SMTP information, administrator email address
  - ILOM configuration. However, ILOM configurations typically remain undamaged even when the server software fails.
- The recovery procedure *does* restore these configuration settings:
  - The network configuration that existed at the end of last successful run of the `/usr/local/bin/ipconf` utility
  - The SSH identities for the server
  - The `root`, `celladmin` and `cellmonitor` users

> **See Also:**
>
> *Oracle Exadata Storage Server Software User's Guide* for information about the `ALTER CELL` command

- The rescue procedure does not examine or reconstruct data disks or data partitions on the system disks. If there is data corruption on the grid disks, then do not use this rescue procedure. Instead, use the rescue procedures for Oracle Database and Oracle ASM.

After a successful rescue, you must reconfigure the server. If you want to preserve the data, then import the cell disks. Otherwise, you must create new cell disks and grid disks.

> **See Also:**
>
> *Oracle Exadata Storage Server Software User's Guide* for information on configuring cells, cell disks, and grid disks using the CellCLI utility

Rescuing a Server Using the CELLBOOT USB Flash Drive

> **Caution:**
>
> Follow the rescue procedure with care to avoid data loss.

To rescue a server using the CELLBOOT USB flash drive:

1. Connect to the Oracle ILOM service processor (SP) of the rescued server. You can use either HTTPS or SSH.
2. Start the server. As soon as you see the splash screen, press any key on the keyboard. The splash screen is visible for only 5 seconds.
3. In the displayed list of boot options, select the last option, `CELL_USB_BOOT_CELLBOOT_usb_in_rescue_mode`, and press Enter.
4. Select the rescue option, and proceed with the rescue.
5. At the end of the first phase of the rescue, choose the option to enter the shell. Do not restart the system
6. Log in to the shell using the rescue `root` password.
7. Use the `reboot` command from the shell.
8. Press F8 as the server restarts and before the splash screen appears. Pressing F8 accesses the boot device selection menu.
9. Select the RAID controller as the boot device. This causes the server to boot from the hard disks.

> **Note:**
>
> Additional options might be available that allow you to enter a rescue mode Linux login shell with limited functionality. Then you can log in to the shell as the `root` user with the password supplied by Oracle Support Services, and manually run additional diagnostics and repairs on the server. For complete details, contact your Oracle Support Services representative.

Reconfiguring the Rescued Storage Server

After a successful rescue, you must configure the server. If the data partitions were preserved, then the cell disks are imported automatically during the rescue procedure.

1. For any replaced servers, re-create the cell disks and grid disks.
2. Log in to the Oracle ASM instance, and set the disks to `ONLINE` using the following command for each disk group:

   ```
   SQL> ALTER DISKGROUP disk_group_name ONLINE DISKS IN FAILGROUP \
   cell_name WAIT;
   ```

3. Reconfigure the cell using the `ALTER CELL` command. The following example shows the most common parameters:

   ```
   CellCLI> ALTER CELL
   smtpServer='my_mail.example.com', -
   smtpFromAddr='john.doe@example.com', -
   smtpFromPwd=email_address_password, -
   smtpToAddr='jane.smith@example.com', -
   notificationPolicy='critical,warning,clear', -
   notificationMethod='mail,snmp'
   ```

4. Re-create the I/O Resource Management (IORM) plan.
5. Re-create the metric thresholds.

> **See Also:**
>
> *Oracle Exadata Storage Server Software User's Guide* for information about IORM plans and metric thresholds

Recreating a Damaged CELLBOOT USB Flash Drive

If the CELLBOOT USB flash drive is lost or damaged, then you can create another one.

To create a CELLBOOT flash drive:

1. Log in to the server as the `root` user.
2. Attach a new USB flash drive with a capacity of 1 to 8 GB.
3. Remove any other USB flash drives from the system.
4. Change directories:

   ```
   cd /opt/oracle.SupportTools
   ```

5. Copy the server software to the flash drive:

   ```
   ./make_cellboot_usb -verbose -force
   ```

13
Reconfiguring the Hardware

This chapter describes how to make basic adjustments. It contains the following topics:

Changing Component Passwords

Oracle recommends that you change the passwords for the various rack components after the initial hardware and software configuration:

> **See Also:**
>
> Default User Accounts for Oracle Zero Data Loss Recovery Appliance for a list of default users and passwords of the various Recovery Appliance components.

Changing the Compute Server Passwords

You can change the user accounts and GRUB passwords on the compute servers. default user accounts on a compute server are `root` and the software owner account. Typically the software owner account is `oracle`.

Changing the User Account Password

A compute server has two default user accounts: `root` and the software owner (typically `oracle`).

To change a user account password on a compute server:

1. Log in to the compute server as root.

2. Use the Linux `passwd` command. This example changes the password of the `oracle` account:

```
# passwd oracle
Changing password for user oracle.
Changing password for oracle
(current) UNIX password: old_password
New UNIX password: new_password
Retype new UNIX password: new_password
```

Changing the GRUB Account Password

To change the GRUB account password on a compute server:

1. Log in to the compute server as `root`.
2. Use the following command:

```
# grub-md5-crypt
```

   The system prompts for a new password twice.

3. Copy the string generated by `grub-md5-crypt` to the copy buffer.
4. Open the `/boot/grub/grub.conf` file in a text editor and locate the password setting. It looks similar to the following:

```
password --md5 hashed_string
```

5. Replace the existing hashed string with the one copied from the `grub-md5-crypt` command output.
6. Save the file.

Changing the Storage Server Passwords

The default user accounts on storage servers are `root`, `celladmin`, and `cellmonitor`.

To change a user account password on a storage server:

1. Log in to the compute server as `root`.
2. Use the Linux `passwd` command. This example changes the password of the `celladmin` account:

```
# passwd celladmin
Changing password for user celladmin.
Changing password for celladmin
(current) UNIX password: old_password
New UNIX password: new_password
Retype new UNIX password: new_password
```

Creating New Users and Changing Passwords on a PDU

The default account user for the power distribution unit (PDU) is `admin`. You can create up to five users and set their passwords.

To create users and passwords on a PDU:

1. Open a web browser and enter the IP address of the PDU in the address field.

The Current Measurement page appears.

2. Click **Network Configuration** in the upper left of the page.
3. Log in as the `admin` user on the PDU metering unit.
4. Locate the Admin/User fields. Only letters and numbers are allowed for user names and passwords.
5. Enter up to five user names and passwords in the Admin/Users fields.
6. Designate each user as either an administrator or a user.
7. Click **Submit** to set the user names and passwords.

Changing the Oracle ILOM Password

The default user account for Oracle Integrated Lights Out Manager (ILOM) is `root`.

To change the password for an Oracle ILOM user:

1. Connect to Oracle ILOM using SSH, and log in as `root`.
2. Use the following command to change the password:

   ```
   set /SP/users/user_name password
   ```

This example changes the password for `user1`:

```
set /SP/users/user1 password

Changing password for user /SP/users/user1/password...
Enter new password:********
Enter new password again:********
New password was successfully set for user /SP/users/user1
```

Changing an InfiniBand Switch Password

The default user accounts on the InfiniBand switches are `root`, `ilom-admin`, `ilom-user`, `ilom-operator`, and `nm2user`.

To change a password for a user on an InfiniBand switch:

1. Connect to the InfiniBand switch using SSH and log in:

   ```
   ssh user_name@switch_name
   ```

   In the preceding command, *user_name* is the name of the user, and *switch_name* is the name of the InfiniBand switch.
2. Check the firmware version of the switch.
3. Use Oracle ILOM to change the password:

   ```
   ssh -l ilom-admin switch_name

   set /SP/users/user_name password
   ```

> **See Also:**
>
> *Sun Datacenter InfiniBand Switch 36 User's Guide* at
>
> http://docs.oracle.com/cd/E36265_01/index.html

Changing the Cisco Ethernet Switch Password

There are no user accounts on the Cisco Ethernet switch. The password, which is called the "enable" password, is not specific to a user account.

To change the password for the Cisco Ethernet switch:

1. Change to enable mode:

```
Switch> enable
```

2. Set the password:

```
rasw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
rasw-ip(config)# enable password password
rasw-ip(config)# enable secret password
rasw-ip(config)# end
rasw-ip# write memory
*Sep 15 14:25:05.893:%SYS-5-CONFIG_I:Configured from console by console
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

3. Save the current configuration:

```
rasw-ip# copy running-config startup-config
```

4. Exit from the session:

```
rasw-ip# exit
```

Changing the Recovery Appliance Database User Password

Oracle recommends that you change the password of the Recovery Appliance default database user account.

To change the Recovery Appliance database user password:

1. Log in to the compute server as `rasys`.
2. Issue the following RACLI command.

```
racli alter rasys_user --password
```

3. When prompted, enter and confirm the new password.

```
Enter rasys Password:
Confirm new rasys Password:
```

The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.

> **See Also:**
>
> Default User Accounts for Oracle Zero Data Loss Recovery Appliance for information about the default Recovery Appliance database accounts and passwords.

Changing OSB Tape Backup Passwords

The OSB tape backup default accounts include the application users `admin` and `oracle`, and the encryption key wallet account.

To change an OSB tape backup application user password:

1. Log in to the compute server as `root`.
2. Use the `obtool chuser` command.

   This example changes the password of the `admin` account:

```
# obtool chuser -q admin
Password:********
Password (again):********
#
```

To change the password for the OSB tape backup encryption key wallet:

1. Log in to the compute server as `root`.
2. Run the following commands:

```
# export OBCONFIG=/osbcat/etc/obconfig
# obcm chpass --keywallet
Old wallet password:********
New wallet password:********
New wallet password (again):********
Wallet password has been reset
#
```

> **See Also:**
>
> Default User Accounts for Oracle Zero Data Loss Recovery Appliance for information about the default OSB tape backup accounts and passwords for Recovery Appliance.

Changing the DNS Servers

You can change the configuration settings for the Domain Name System (DNS) servers after the initial setup. All servers and switches in Recovery Appliance must reference the same DNS servers. This section explains how to set the servers and switches to the same DNS servers. Oracle recommends changing the servers one at a time.

- "Changing the DNS Server Address on an InfiniBand Switch"
- "Changing the DNS Server Address on the Cisco Ethernet Switch"

- "Changing the DNS Server Address on the Compute Server"
- "Changing the DNS Server on Storage Servers"

Changing the DNS Server Address on an InfiniBand Switch

To change the DNS server address on an InfiniBand switch, connect to Oracle ILOM and log in as the `ilom-admin` user. You can choose from one of the following methods:

- **Oracle ILOM web interface**

  Select the **Configuration** tab and set the DNS server addresses.

- **Oracle ILOM CLI**

  Use the following command to set the DNS:

  ```
  set /SP/clients/dns nameserver=dns_ip
  ```

  In the preceding command, *dns_ip* is the IP address of the DNS server. For multiple servers, enter a comma-separated list. For example:

  ```
  set /SP/clients/dns nameserver=dns_ip1,dns_ip2,dns_ip3
  ```

Changing the DNS Server Address on the Cisco Ethernet Switch

To change the DNS server address on the Cisco Ethernet switch:

1. Access the Ethernet switch using SSH, and log in as the `admin` user with the administrator password.

   If SSH is not configured, then use Telnet to access the switch.

2. Change to enable mode:

   ```
   Switch> enable
   ```

3. List the DNS IP addresses currently in use:

   ```
   Switch# show running-config
   ```

4. Erase the current DNS server addresses and any invalid IP addresses. This example erases three IP addresses:

   ```
   Switch# configure terminal
   Enter configuration commands,one per line. End with CNTL/Z.
   Switch(config)# no ip name-server 10.7.7.2
   Switch(config)# no ip name-server 129.148.5.4
   Switch(config)# no ip name-server 10.8.160.2
   Switch(config)# end
   Switch# write memory
   Building configuration...
   Compressed configuration from 2603 bytes to 1158 bytes [OK ]
   ```

5. Configure up to three DNS servers. The following example configures three new DNS servers:

   ```
   Switch# configure terminal
   Enter configuration commands,one per line. End with CNTL/Z.
   Switch(config)# ip name-server 10.7.7.3
   Switch(config)# ip name-server 129.148.5.5
   Switch(config)# ip name-server 10.8.160.1
   Switch(config)# write memory
   ```

```
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

6. Verify the changes:

```
Switch# show running-config
Building configuration...
...
ip domain-name example.com
ip name-server 192.168.10.2
ip name-server 192.168.10.3
ip name-server 192.168.10.4
...
```

7. Save the running configuration as the startup configuration:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 14343 bytes to 3986 bytes[OK]
```

8. End the session:

```
Switch# exit
```

Changing the DNS Server Address on the Compute Server

To change the DNS server address on the compute servers:

1. Log in as the `root` user.
2. Open `/etc/resolv.conf` file in a text editor.
3. Locate the name server setting for the DNS server and change it to the new value. For example:

```
search        example.com
nameserver    10.7.7.3
```

4. Save the file and close the editor.
5. Set the DNS server in the server ILOM:

```
ipmitool sunoem cli 'set /SP/clients/dns nameserver=dns_ip'
```

In the preceding command, *dns_ip* is the IP address of the DNS server. For multiple DNS servers, enter a comma-separated list. For example:

```
set /SP/clients/dns nameserver=dns_ip1,dns_ip2,dns_ip3
```

6. Repeat this procedure for each compute server.

Changing the DNS Server on Storage Servers

To change the DNS server or NTP settings on a storage server:

1. Log in to the storage server as the `root` user.
2. Follow steps 1 through 5 of My Oracle Support Doc ID 1188080.1 to take the server offline.
3. Shut down the cell services and `ocrvottargetd` service:

```
# cellcli -e alter cell shutdown services all
# service ocrvottargetd stop
```

4. Use the `ipconf` utility to change the DNS settings:

   ```
   # /usr/local/bin/ipconf
   ```

5. Restart the server:

   ```
   # shutdown -r now
   ```

6. Follow steps 7 through 9 of My Oracle Support note 1188080.1 to put the server back online.
7. Repeat this procedure for each storage server.

Changing the NTP Servers

You can change the configuration settings for the Network Time Protocol (NTP) servers after the initial setup. All servers and switches in Recovery Appliance must reference the same NTP servers, so that they are synchronized to the same time. This section explains how to set the Recovery Appliance servers and switches to the same NTP server addresses. Oracle recommends changing each server one at a time.

This section contains the following tasks:

- "Setting the NTP Server Address on the Compute Servers"
- "Setting the NTP Server Address on the InfiniBand Switches"
- "Set the NTP Server Address on the Cisco Ethernet Switch"
- "Setting the NTP Server on the Storage Servers"
- "Restart Recovery Appliance"

> **Note:**
>
> - You must restart Recovery Appliance after completing these tasks.
> - You can configure one or two NTP servers for use with Recovery Appliance.

Setting the NTP Server Address on the Compute Servers

To set the NTP server address on a compute server:

1. Log in to a compute server as `root`.
2. Stop the NTP services on the server:

   ```
   service ntpd stop
   ```

3. Open `/etc/ntp.conf` in a text editor and enter the IP address of the new NTP server.
4. Restart the NTP services on the compute server:

   ```
   service ntpd start
   ```

5. Repeat this task for the second compute server.

Setting the NTP Server Address on the InfiniBand Switches

> **Note:**
>
> Do not manually edit the files on the InfiniBand switches.

To set the NTP server address on a Sun Datacenter InfiniBand Switch 36:

1. Log in to the switch as the `ilom-admin` user.
2. Set the date, time zone, and Network Time Protocol (NTP), using one of the following methods:
   - The Configuration page of the ILOM graphical interface.
   - The following CLI commands:

     ```
     set /SP/clock timezone=preferred_tz
     set /SP/clients/ntp/server/1 address=ntp_ip1
     set /SP/clients/ntp/server/2 address=ntp_ip2
     set /SP/clock usentpserver=enabled
     ```

     In the preceding commands, *preferred_tz* is the preferred time zone, and *ntp_ip1* and *ntp_ip2* are the NTP server IP addresses. You must configure at least one NTP server; the second one is optional.
3. Repeat this task for the other InfiniBand switches.

Set the NTP Server Address on the Cisco Ethernet Switch

To set the NTP server on the Cisco Ethernet switch:

1. Open a connection to the Ethernet switch using SSH, and log in as `admin`. If SSH is not configured, then use Telnet.
2. Change to enable mode.

   ```
   Switch> enable
   ```
3. Review the current configuration:

   ```
   Switch# show running-config
   ```
4. Erase the current NTP server configuration and any invalid NTP addresses. In the following example, the current IP addresses are 10.10.10.1 and 10.8.8.1.

   ```
   Switch# configure terminal
   Enter configuration commands,one per line. End with CNTL/Z.
   Switch(config)# no ntp server 10.10.10.1
   Switch(config)# no ntp server 10.8.8.1
   Switch(config)# end
   Switch# write memory
   Building configuration...
   Compressed configuration from 2603 bytes to 1158 bytes [OK ]
   ```
5. Configure up to two NTP servers. The following example uses 10.7.7.1 and 10.9.9.1 as the new IP addresses:

```
Switch# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
Switch(config)# ntp server 10.7.7.1 prefer
Switch(config)# ntp server 10.9.9.1
Switch(config)# end
Switch# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

6. Verify the changes:

```
Switch# show running-config

Building configuration...
...
ntp server 192.168.10.10 prefer
...
```

7. Save the configuration:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 14343 bytes to 3986 bytes[OK]
```

8. Exit from the session:

```
Switch# exit
```

Setting the NTP Server on the Storage Servers

To set the NTP server on a storage server:

1. Log in to the storage server as the `root` user.
2. Follow steps 1 through 5 of My Oracle Support note 1188080.1, to take the server offline.
3. Shut down the cell services and `ocrvottargetd` service:

```
# cellcli -e alter cell shutdown services all
# service ocrvottargetd stop
```

4. Use the `ipconf` utility to change the NTP settings:

```
# /usr/local/bin/ipconf
```

5. Restart the server:

```
# shutdown -r now
```

6. Follow steps 7 through 9 of My Oracle Support note 1188080.1, to bring the server back online.

Restart Recovery Appliance

After changing the servers and switches, restart Recovery Appliance.

> **See Also:**
>
> "Powering On and Off a Recovery Appliance Rack"

Changing the Time Zone Settings

This section provides information about changing the time zones on Recovery Appliance after the initial configuration and deployment. You must modify the following components when changing the time zone settings:

- Storage servers
- Compute servers
- Sun Datacenter InfiniBand Switch 36 switches
- Cisco switch

> **Note:**
>
> Before changing the time zone settings, you must stop the cell services on the storage servers and the Oracle Clusterware Services.

The following tasks describe how to change the time zone settings on the components:

- "Changing the Time Zone Settings on the Storage Servers"
- "Change Time Zone Settings on the Compute Servers"
- "Change the Time Zone Settings on the InfiniBand Switches"
- "Change the Time Zone Settings on the Cisco Switch"

Changing the Time Zone Settings on the Storage Servers

Complete the setting changes to all storage servers before changing the settings on the compute servers.

To change the time zone setting on a storage server:

1. Log in as the `root` user on the storage server.
2. Stop the processes on the cells:

   ```
   # cellcli -e alter cell shutdown services all
   ```

3. Run the `ipconf` script:

   ```
   # /opt/oracle.cellos/ipconf
   ```

4. Proceed through the script prompts until you see the time zone prompts. Do not change any other settings.

   The following example changes the time zone from McMurdo, Antarctica, to New York City, USA. The number for the United States is 230.

   ```
   The current timezone: Antarctica/McMurdo
   Do you want to change it (y/n) [n]: y

   Setting up local time...

   1) Andorra
   2) United Arab Emirates
   ```

```
3) Afghanistan
.
.
.
15) Aruba
16) Aaland Islands
Select country by number, [n]ext, [l]ast: 230

Selected country: United States (US). Now choose a zone

1) America/New_York
2) America/Detroit
3) America/Kentucky/Louisville
.
.
.
15) America/North_Dakota/New_Salem
16) America/Denver
Select zone by number, [n]ext: 1

Selected timezone: America/New_York
Is this correct (y/n) [y]:
```

5. Proceed through the remaining prompts, but do not change any other values.
6. Ensure that the time zone changes appear in the following files. Examples of the changes are shown for the files.
   - `/opt/oracle.cellos/cell.conf`

     ```
     $VAR1 = {
               'Hostname' => 'xdserver.us.example.com',
               'Ntp servers' => [
                                  '10.141.138.1'
                                ],
               'Timezone' => 'America/New_York',
     ```

   - `/etc/sysconfig/clock`

     ```
     ZONE="America/New_York"
     UTC=false
     ARC=false
     #ZONE="Antarctica/McMurdo"
     #ZONE="America/New_York"
     #ZONE="America/Los_Angeles"
     ```

     The uncommented value (not preceded by #) is the current setting.

   - `/etc/localtime`

     Use the `strings /etc/localtime` command to see the printable characters in the file. The last line includes the time zone setting.

     ```
     EST5EDT,M3.2.0,M11.1.0
     ```

7. Restart the server.
8. Use the `date` command to see the current time zone. The following is an example of the output from the command:

   ```
   # date
   Thu Jun 12 16:43:47 EDT 2014
   ```

9. Review the `$ADR_BASE/diag/asm/cell/`*`host_name`*`/alert.log` file. Verify that the time that you restarted the processes matches the current and correct time.

Change Time Zone Settings on the Compute Servers

To change the time zone setting on a compute server:

1. Log in as the `root` user on the compute server.
2. Shut down Oracle Clusterware Services:

   ```
   # GI/bin/crsctl stop crs
   ```

3. Copy the `clock` file from one of the storage servers:

   ```
   # scp root@storage_cell:/etc/sysconfig/clock /etc/sysconfig/clock
   ```

4. Disable Cluster Ready Services (CRS):

   ```
   # GI/bin/crsctl disable crs
   ```

   This step prevents Cluster Ready Services (CRS) from restarting automatically when you restart the server.
5. Restart the compute server.
6. Use the `date` command verify the change for the time zone.
7. Enable CRS:

   ```
   # GI/bin/crsctl enable crs
   ```

   This step restores CRS to its previous configuration setting.
8. Start CRS:

   ```
   # GI/bin/crsctl start crs
   ```

Change the Time Zone Settings on the InfiniBand Switches

The following procedure describes how to change the time zone setting on a Sun Datacenter InfiniBand Switch 36:

1. Connect to the switch using SSH, and log in as the `root` user.
2. Log in to Oracle ILOM using the web address:

   `http://switch_alias`
3. Select the **Configuration** tab.
4. Select the **Clock** tab.
5. Ensure the Synchronize Time Using NTP field is enabled.
6. Enter the correct IP address for the servers.
7. Click **Save**.

Change the Time Zone Settings on the Cisco Switch

To change the time zone setting on the Cisco switch:

1. Use Telnet to connect to the Cisco switch as the `root` user.
2. Use the `enable` command to enter privileged mode.

3. Use the `configure terminal` command to begin configuration.
4. Set the clock using the following command:

   ```
   clock timezone zone hours_offset [minutes_offset]
   ```

   In the preceding command, *zone* is the name of the time zone to be displayed when standard time is in effect, *hours_offset* is the hours offset from UTC, and *minutes_offset* are the minutes offset from UTC. The default time zone is UTC.

   To set summer time (daylight savings time), use the following command:

   ```
   clock summer-time zone recurring [week day month hh:mm week day month   \
         hh:mm [offset]]
   ```

   In the preceding command, *week* is the week of the month, between 1 to 5, *day* is the day of the week, such as Sunday or Monday, *month* is the month, such as January or June, *hh:mm* is the time in 24-hour format, and offset is the number of minutes to add during summer time. The default for `hh:mm` is 60. Summer time is disabled by default.

   The following is an example of setting the time zone to US Eastern time with summer time enabled:

   ```
   $ telnet raeth-ip
   Connected to raeth-ip
   Escape Character is 'CTRL+]'

   User Access Verification

   Password: password
   raeth-ip> enable
   Password: password
   raeth-ip# configure terminal
   Enter configuration commands, one per line. End with CNTL/Z.
   raeth-ip(config)# clock timezone EST -5
   raeth-ip(config)# clock summer-time EDT recurring
   raeth-ip(config)# end
   raeth-ip# write memory
   Building configuration...
   Compressed configuration from 6421 bytes to 2041 bytes[OK]
   raeth-ip# show clock
   12:03:43.516 EDT Wed May 12 2012
   ```

Part IV
Appendixes

Zero Data Loss Recovery Appliance Owner's Guide appendixes are:

- RACLI Overview
- Default User Accounts for Oracle Zero Data Loss Recovery Appliance
- Error Messages
- Replacement Units
- In-Rack Cabling Tables for Recovery Appliance X9M
- In-Rack Cabling Tables for Recovery Appliance X8M
- In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5
- In-Rack Cabling Tables for Recovery Appliance X4

A

RACLI Overview

This topic provides a summary of all the RACLI commands.

NAME

racli - Recovery Appliance Command Line Interface

DESCRIPTION

The racli is a simplified and controlled environment for users to manage their Recovery Appliance via the command line.

SYNOPSIS

racli <COMMAND> <OBJECT>

HELP

Displays information about racli commands and their objects.

Usage: racli help [<COMMAND> [<OBJECT>]]

Options:

`racli help`
General usage information for the racli help command.

`racli help <COMMAND>`
Detailed information specific to the COMMAND requested.

`racli help <COMMAND> <OBJECT>`
Specific detailed information on the COMMAND and OBJECT relationship.

Command	Description
racli add	Adds an object to the Recovery Appliance.
racli add acfs	Specifies that the Automatic Storage Management Cluster (ACFS) should be added to the Recovery Appliance schema.
racli add cloud_key	Generates the public and private key files required for OCI access.
racli add cloud_location	Adds a cloud storage object to the Recovery Appliance.
racli add cloud_user	Adds a cloud user to the Recovery Appliance schema.
racli add credential	Adds a user credential to the wallet of the Recovery Appliance.
racli add credential_wallet	Adds a wallet object to the Recovery Appliance.
racli add host	Adds a host to those available on the Recovery Appliance.

Command	Description
racli add keystore	Adds a keystore to the Recovery Appliance.
racli add nfs	Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.
racli add rasys_wallet	Creates the rasys user wallet.
racli add storage_cells	Prepares new Recovery Applicance storage cells for storage location expansion.
racli add storage_disks	Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.
racli add tag	Specifies the diagnostic TAG to add to all nodes of the Recovery Appliance**Usage:** racli add tag --tag=myTagName --level=myLevel**Options:**Name of the new tag to be added.
racli add vpc_user	Creates a new Virtual Private Catalog (VPC) user account in the Recovery Appliance schema.
racli alter	Modifies the properties on objects supported in the Recovery Appliance schema.
racli alter check	Modifies a check on the Recovery Appliance.
racli alter cloud_key	Modifies the information associated with a cloud key.
racli alter cloud_location	Modifies the properties of cloud storage supported on the Recovery Appliance.
racli alter cloud_user	Modifies information about a cloud user.
racli alter credential_wallet	Changes the password on a wallet object in the Recovery Appliance.
racli alter keystore	Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.
racli alter parameter	Modifies the specified parameter whether it should be included or ignored as a check.
racli alter passphrase	Modifies password used to access the Recovery Appliance.
racli alter rasys_user	Modifies the Recovery Appliance rasys user account password and management of the password.
racli alter vpc_user	Modifies an existing Recovery Appliance user account and changes its password.
racli configure	Configure an object from operations performed by the Recovery Appliance.
racli configure archive_group	Configures archive log grouping.
racli disable	Disables an object from operations performed by the Recovery Appliance.
racli disable aide	Disables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.
racli disable check	Disables a specific check object from check run operations performed on the Recovery Appliance.
racli disable cloud_location	Disables a cloud storage object for operations performed by the Recovery Appliance.
racli disable keystore	Disables the keystore.

Command	Description
racli disable nfs	Disables the specified Network File System (NFS).
racli enable	Enables an object for operations performed by the Recovery Appliance.
racli enable acfs	Enables the specified Automatic Storage Management Cluster (ACFS).
racli enable aide	Enables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.
racli enable check	Enables a check on the Recovery Appliance.
racli enable cloud_location	Enables a cloud storage object for operations performed by the Recovery Appliance.
racli enable keystore	Opens the keystore so that it can be used.
racli enable nfs	Enables the specified Network File System (NFS).
racli export	Used for exporting data.
racli export appliance	Creates an export bundle for the Recovery Appliance.
racli install	Installs software.
racli install appliance	Installs software to the Recovery Appliance software.
racli install okv_endpoint	Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance.
racli list	Lists information about objects given in the Recovery Appliance schema.
racli list check	List checks available on the Recovery Appliance.
racli list cloud_key	List details in short (default) or long format about a specified cloud key or all cloud keys.
racli list cloud_location	Lists details about the cloud storage known to the Recovery Appliance.
racli list cloud_user	List details about the specified cloud user, or all cloud users.
racli list compute_nodes	Lists the names of the compute nodes available on the Recovery Appliance.
racli list grid_home	Lists details about grid_home.
racli list host	Lists details about all hosts available to the Recovery Appliance.
racli list nfs	Lists details about one or more Network File System mounts.
racli list okv_endpoint	Lists details about the OKV end points known to the Recovery Appliance.
racli list oracle_home	Lists details about oracle_home.
racli list parameter	Lists the names and value of the init parameters for the Recovery Appliance.
racli list storage_cells	Lists names of storage cells available to the Recovery Appliance.
racli list tag	List diagnostic tags available on the Recovery Appliance.
racli patch	The action to patch a specified object.
racli patch appliance	Patch the Recovery Appliance Software.
racli remove	Removes objects from the Recovery Appliance schema.

Command	Description
racli remove acfs	Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.
racli remove cloud_key	Removes the specified cloud key, but only if no cloud users are associated with the key.
racli remove cloud_location	Removes the specified cloud storage from the Recovery Appliance, unless the cloud archive job template is tied to the library.
racli remove cloud_user	Removes a cloud user.
racli remove grid_home	Removes the specified grid home.
racli remove host	Removes the specified host.
racli remove keystore	Removes the specified keystore from the Recovery Appliance schema.
racli remove nfs	Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.
racli remove oracle_home	Removes the specified oracle home.
racli remove osb_pieces	Removes all pieces that are no longer in the ZDLRA catalog from the OSB catalog.
racli remove tag	Removes the specified diagnostic TAG from all nodes of the Recovery Appliance.
racli remove vpc_user	Removes the specified user account from the Recovery Appliance metadata schema.
racli run	Performs an action for an object to the Recovery Appliance.
racli run check	Runs the enabled checks for the Recovery Appliance.
racli run diagnostics	Runs and collects the required diagnostics on the Recovery Appliance.
racli start	The action to start a specified service (object).
racli start nfs	Start NFS services of the Recovery Appliance.
racli status	Check status of objects associated with the Recovery Appliance.
racli status aide	Returns the status of the Advanced Intrusion Detection Environment (AIDE).
racli status appliance	Returns a summary of the status of the appliance.
racli status cloud_location	Lists the status of a cloud storage object for operations performed by the Recovery Appliance.
racli status disk_group	Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.
racli status keystore	Returns the status of the keystore in use by the Recovery Appliance.
racli status nfs	Returns the status of the Network File Service (NFS).
racli status okv_endpoint	Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.
racli stop	The action to stop a specified service (object).
racli stop nfs	Stop NFS services of the Recovery Appliance.
racli uninstall	Backs out a portion of the installation.

Command	Description
racli uninstall appliance	Backout a step during install.
racli uninstall okv_endpoint	Uninstalls the specified OKV end point.
racli update	Update a component of the Recovery Appliance.
racli update aide	Updates the Advanced Intrusion Detection Environment (AIDE).
racli update parameter	Update init parameters in the Recovery Appliance.
racli update storage_location	Update RA pool data file with the most current number of storage cells.
racli update support_tools	Updates the AHF and Exachk versions using the staged binaries zip files.
racli upgrade	The action to upgrade a specified object to the next version.
racli upgrade appliance	Upgrade the Recovery Appliance Software Versions.
racli version	Provides version information for software running on the Recovery Appliance.

racli add

Adds an object to the Recovery Appliance.

Options

Details

Adds an object to the Recovery Appliance.

racli add acfs

Specifies that the Automatic Storage Management Cluster (ACFS) should be added to the Recovery Appliance schema.

Syntax

```
racli add acfs --name=myAcfsName --disk_group=myDGName --
volume_name=myVolumeName --mount_point=myPath --size=SIZE
```

Options

`--name`
Name to be given to the ACFS mount object

`--disk_group`
Name of the ASM disk group to use for this ACFS mount.

`--volume_name`
Name of the ASM VOLUME to mount the ACFS on.

`--mount_point`
Path to mount the ACFS object.

--size
Size of the ACFS object.

Details

Specifies that the Automatic Storage Management Cluster (ACFS) should be added to the Recovery Appliance schema.

racli add cloud_key

Generates the public and private key files required for OCI access.

Syntax

```
racli add cloud_key --key_name=KEY_NAME [--fingerprint=PUBFINGERPRINT --
pvt_key_path=PVTKEYFILE]
```

Options

--key_name
The name of the `cloud_key` object to create.

--fingerprint
Specifies the name of the fingerprint for the public key that is paired with the specified private key. The fingerprint is used to authenticate the request.

--pvt_key_path
The name of a private key used to sign BMC API requests. The key file must be in PEM format.

Details

Generates the public and private key files required for OCI access. It also associates fingerprint and pre-existing private key files.

The first syntax form generates the public and private keys required for OCI access. The user takes the generated public key to get fingerprint from OCI console.

```
racli add cloud_key --key_name=KEY_NAME
```

The second syntax form associates fingerprint and private key files that already exist.

```
racli add cloud_key --key_name=KEY_NAME --fingerprint=PUBFINGERPRINT --
pvt_key_path=PVTKEYFILE
```

racli add cloud_location

Adds a cloud storage object to the Recovery Appliance.

Syntax

```
racli add cloud_location --cloud_user=CLOUD_USER_NAME --host=HOST_URL --
bucket=OCI_BUCKET_NAME [--enable_archive=TRUE] [--
archive_after_backup=NUMBER:{DAYS|YEARS} --streams=NUMBER --
proxy_host=HTTP_SERVER --proxy_port=HTTP_PORT --proxy_id=HTTP_USER --
```

```
proxy_pass=HTTP_PASS --import_all_trustcert=X509_CERT_PATH --
retain_after_restore=NUMBER:HOURS]
```

Options

`--archive_after_backup`
The number in days or years for period of time before backups are moved to archive state. It takes the form "--archive_after_backup=NUMBER:DAYS" or "--archive_after_backup=NUMBER:YEARS", where NUMBER is a positive integer. If not specified, the default is 2 days.

`--bucket`
The name of the bucket where the backup will go. Note that the install tool will create the specified bucket if it does not exist.

`--cloud_user`
Previously configured cloud_user object with all authentication requirements.

`--enable_archive`
Enables archival of backup data after the time interval specified by the `--archive_after_backup` parameter. If not specified, by default the cloud archival functionality will be disabled, which is equivalent to a cloud bucket that is just cold storage.
If cloud location is enabled for archival, use "`racli disable cloud_location`" and "`racli enable cloud_location`" later as pause/resume functionality.

`--host`
Host name for the Oracle Bare Metal Cloud account.

`--import_all_trustcert`
Import all X509 certificates from Java truststore.

`--streams`
The number of streams to be associated with this specific cloud_location. Directly translates to the number of sbt_attrib sets created.

`--proxy_host`
HTTP proxy server.

`--proxy_id`
HTTP proxy server username. Used for OCI and stored as part of the `alias_http_proxy` in location wallet.

`--proxy_pass`
HTTP proxy server password. Stored as part of the `alias_http_proxy` in location wallet.

`--proxy_port`
HTTP proxy server port. Used for OCI-C support.

`--retain_after_restore`
Specifies the number of hours to retain information after a restore.
It takes the form "--retain_after_restore=NUMBER:HOURS". It has a maximum value of "240:HOURS".

Details

Adds a cloud storage object to the Recovery Appliance. This is equivalent to the `sbt_library` used for copy-to-tape but is specific to copy-to-cloud. The user's cloud credentials are required.

Examples

This example adds a new cloud storage object to the Recovery Appliance schema.

```
racli add cloud_location --cloud_user=samu ---host=https://cloud_uri/oci -
bucket=sam_backup
```

racli add cloud_user

Adds a cloud user to the Recovery Appliance schema.

Syntax

```
racli add cloud_user --user_name=CLOUD_USER_NAME --key_name=CLOUD_KEY_NAME
--user_ocid=UOCID --tenancy_ocid=TOCID [--compartment_ocid=COCID]
```

Options

`--user_name`
The name to be associated with this particular cloud user.

`--key_name`
The specific cloud key to be associated with this cloud user.

`--tenancy_ocid`
The tenancy OCID for the Oracle Bare Metal Cloud account.

`--user_ocid`
The user OCID for the Oracle Bare Metal Cloud account.

`--compartment_ocid`
The compartment OCID within the tenancy of the Oracle Bare Metal Cloud Account.

Details

Adds a cloud user to the Recovery Appliance schema.

racli add credential

Adds a user credential to the wallet of the Recovery Appliance.

Syntax

```
racli add credential --user_name=OS_USER_ID --alias=ALIAS_ID --
wallet_path=PATH_TO_WALLET
```

Options

`--user_name`
The OS user name associated with the credential to be added.

`--alias`
The alias you want to associate to the OS username that corresponds to the credential to be added.

`--wallet_path`
The path to the wallet where you want to store the credential.

Details

Adds a user credential to the wallet of the Recovery Appliance.

Examples

This example adds a credetinal for OS user *charlie* with an alias of *sierra* into the designated wallet.

```
racli add credential --wallet_path=/raacfs/raadmin/cloud/test --
user_name=charlie --alias=sierra
```

When prompted, enter the password for the OS user for which the credential is to be added.

```
Enter Password [charlie]:
Confirm Password [charlie]:
```

When prompted, enter the password for the wallet the credential is to be added.

```
Enter Wallet Password:
Confirm Wallet Password:
```

racli add credential_wallet

Adds a wallet object to the Recovery Appliance.

Syntax

```
racli add credential_wallet --wallet_path=PATH_TO_WALLET
```

Options

`--wallet_path`
Path the credential wallet will reside.

Details

Adds a wallet object to the Recovery Appliance. You are prompted to enter new passwords for the keystore and then the wallet. The credentials to access the Recovery Appliance encryption keystore are saved in this wallet.

If no wallet path is specified, the default path is used `/raacfs/raadmin/ra_wallet/`.

Examples

This example adds a credential wallet in the defualt path.

```
racli add credential_wallet
```

When prompted, enter the password for the new credential wallet.

```
Enter Wallet Password:
Confirm Wallet Password:
```

racli add host

Adds a host to those available on the Recovery Appliance.

Syntax

```
racli add host --host_name=HOSTNAME --ip=IPADDRESS --
domain_name=DOMAINNAME
```

Options

`--host_name`
Indicates the name of the host to be added.

`--ip`
Indicates the IP address for the host being added.

`--domain_name`
Indicates the domain name for the host being added.

Details

Adds a host to those available on the Recovery Appliance.

racli add keystore

Adds a keystore to the Recovery Appliance.

Syntax

```
racli add keystore --type=[HSM|SW] --wrl=PATH_TO_WALLET
```

Options

`--type`
Specifies the type of keystore to be created. If not specified, default is HSM.

`--wrl`
Specifies the path to the wallet that will be used for the keystore.

Details

Adds a keystore to the Recovery Appliance. The keystore contains one or more TDE master keys for each Recovery Appliance client database, plus the Recovery Appliance's TDE master key. The per-client TDE master keys are used to encrypt backup pieces that are copied to the cloud.

ATTENTION: The Recovery Appliance database is restarted to activate the keystore; plan for a short outage.

Examples

This example adds a keystore of type HSM to the Recovery Appliance.

```
racli add keystore --type=hsm
```

racli add nfs

Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.

Syntax

```
racli add nfs --nfs_server=myServerName --nfs_export=myExportPath --
name=myNfsMountName --mount_options='NFS_OPTIONS'
```

Options

`--name`
Name to be given to the NFS mount object.

`--nfs_server`
Server name where the `myExportPath` resides.

`--nfs_export`
Path on the `myServerName` intending to mount.

`--mount_options`
NFS mounting options. For example: 'rw,tcp,vers=3,timeo=600,actimeo=0'

Details

Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.

racli add rasys_wallet

Creates the rasys user wallet.

Syntax

```
racli add rasys_wallet
```

Options

Details

Creates the rasys user wallet.

racli add storage_cells

Prepares new Recovery Applicance storage cells for storage location expansion.

Syntax

```
racli add storage_cells --cells=cell1,cell2,...celln
```

Options

`--cells`
The comma-separated list of storage cell names to prepare.

Details

Prepares new Recovery Applicance storage cells for storage location expansion.

- The new storage cells must be the same Exadata image version as the existing storage cells.
- The new storage cells must not have any associated databases. To verify if this is the case, run:

  ```
  cellcli -e list database
  ```

Examples

This example adds the storage cells named *cellhost1* and *cellhost2* to the Recovery Appliance.

```
racli add storage_cells --cells=cell1,cell,...celln2
```

racli add storage_disks

Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.

Syntax

```
racli add storage_disks --cells=cell1,cell2,...celln --priority=high
```

Options

`--cell`
The comma-separated list of names of the newly added storage cells for the new disk groups.

`--priority`
Indicates the priority level for the rebalance. The default is low. Low = 4. High = 32.
This command should only be ran after rebalance is complete.

Details

Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.

You must have previously run `racli add storage_cells` to prepare the cells before adding the storage cells to the existing disk groups.

Run this command only if prompted to do so while running the `racli add storage_cells` command.

You can execute this command while the system is running, but the operation can take longer to complete and can negatively impact performance. Avoid running other I/O intensive tasks such as patching or maintenance during this time.

Examples

The following example adds storage cells named *cellhost1* and *cellhost2* to the storage disks.

```
racli add storage_disks --cells=cell1,cell2,...celln
```

racli add tag

Specifies the diagnostic TAG to add to all nodes of the Recovery Appliance**Usage:** racli add tag --tag=myTagName --level=myLevel**Options:**Name of the new tag to be added.

Syntax

```
racli add tag --tag=myTagName --level=myLevel
```

Options

`--tag`
Name of the new tag to be added.

`--level`
The level of the ANDed diagnostics to include in the tag, available levels are:
level 1 System activity report.
level 2 RA dump txt file.
level 4 RA dump datapump.
level 8 ORDERING_WAIT analysis.
level 16 Corruption.
level 32 Gather racli logs from each node.
level 64 Get RA version information from each node.
level 128 Get store and config files from each node.
level 256 Get root and oracle history from each node.
level 512 Get network information from each node.
level 1024 Run obtool gathering useful information.
level 2048 Get osb logs from each node.
level 4096 Get /usr/etc/ob/xcr from each node.
level 8192 Get backup/admin/state from each node.

level 16384 Get all nodes ASM and ZDLRA alert.logs.
The levels can be ANDed together to create a collection.

Details

Specifies the diagnostic TAG to add to all nodes of the Recovery Appliance

Examples

This example adds a new TAG to collect racli logs and RA version information from each node:

```
racli add tag --tag=new_tag --level=196608
```

racli add vpc_user

Creates a new Virtual Private Catalog (VPC) user account in the Recovery Appliance schema.

Syntax

```
racli add vpc_user --user_name=myVpcUserName
```

Options

`user_name`
The name of the Recovery Appliance catalog user account to add.

Details

Creates a new Virtual Private Catalog (VPC) user account in the Recovery Appliance schema.

Each Recovery Appliance user account is the owner of the virtual private catalog.

- The password is not provided on the command line for security reasons. When prompted, enter a password for the Recovery Appliance user account. The password must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.
- The Recovery Appliance user account being added must not already exist in the metadata database.
- The password on the `DEFAULT` database profile will expire. To avoid the password expiring, change the profile associated with the VPC user.

Because strong password authentication is used, simple passwords can result in ORA-280003 or ORA-20006 errors.

When a `vpc_user` is created, the account is locked. The `DEFAULT` database profile has the following limits.

- Failed attemps is set to 10.
- Password life time is 180 days.
- Password lock time is 1 day.
- Password grace time is 7 days.

Examples

This example adds the user name *ravpc1* to the Recovery Appliance schema.

```
racli add vpc_user --user_name=ravpc1
```

When prompted, enter the password for *ravpc1*.

```
ravpc1 New Password:
```

racli alter

Modifies the properties on objects supported in the Recovery Appliance schema.

Options

Details

Modifies the properties on objects supported in the Recovery Appliance schema.

racli alter check

Modifies a check on the Recovery Appliance.

Syntax

```
racli alter check { --check_name=CHECK_NAME };
```

Options

`--check_name`
Indicates the name of the check to be modified from the Recovery Appliance.

Details

Modifies a check on the Recovery Appliance.

racli alter cloud_key

Modifies the information associated with a cloud key.

Syntax

```
racli alter cloud_key --key_name=KEY_NAME [--fingerprint=PUBFINGERPRINT --
pvt_key_path=PVTKEYFILE]
```

Options

`--key_name`
The name of the cloud key object.

--fingerprint
Specifies the name of the fingerprint for the public key that is paired with the specified private key. The fingerprint is used to authenticate the request.

--pvt_key_path
The name of a private key used to sign requests. The key file must be in PEM format.

Details

Modifies the information associated with a cloud key.

A cloud key has to have a fingerprint associated with it in order to later be associated with a user.

racli alter cloud_location

Modifies the properties of cloud storage supported on the Recovery Appliance.

Syntax

```
racli alter cloud_location --location_name=CLOUD_LOCATION_NAME [--
archive_after_backup=NUMBER:{DAYS|YEARS}] [--streams=NUMBER] [--
proxy_host=HTTP_SERVER] [--proxy_id=HTTP_USER] [--proxy_pass=HTTP_PASS]
[--proxy_port=HTTP_PORT] [--retain_after_restore=NUMBER:HOURS]
```

Options

--archive_after_backup
The number in days or years for period of time before backups are moved to archive state. It takes the form "--archive_after_backup=NUMBER:DAYS" or "--archive_after_backup=NUMBER:YEARS", where NUMBER is a positive integer. If not specified, the default is 2 days.

--location_name
The name of the cloud location where the backup will go. The cloud location correlates to OCI object store bucket. Note that the install tool will create the specified bucket if it does not exist.

--proxy_host
HTTP proxy host.

--proxy_id
HTTP proxy server username. Used for OCI and stored as part of the `alias_http_proxy` in location wallet.

--proxy_pass
HTTP proxy server password. Stored as part of the `alias_http_proxy` in location wallet.

--proxy_port
HTTP proxy server port. Used for OCI-C support.

--streams
The number of streams to be associated with this specific cloud_location. Directly translates to the number of sbt_attrib sets created.
Integer that specifies number of streams to associate with this specific cloud_location.

--retain_after_restore
Specifies the number of hours to retain information after a restore. It takes the form "--retain_after_restore=NUMBER:HOURS". It has a maximum value of "240:HOURS".

Details

Modifies the properties of cloud storage supported on the Recovery Appliance. This is equivalent to the sbt_library used for copy-to-tape. This sbt_library is specific to archive-to-cloud.

racli alter cloud_user

Modifies information about a cloud user.

Syntax

```
racli alter cloud_user --cloud_user=CLOUD_USER_NAME --
key_name=CLOUD_KEY_NAME --user_ocid=UOCID --tenancy_ocid=TOCID --
compartment_ocid=COCID
```

Options

--user_name
The name user wants associated with this particular cloud user.

--key_name
The specific cloud key to be associated with this cloud user.

--tenancy_ocid
The tenancy OCID for the Oracle Bare Metal Cloud account.

--user_ocid
The user OCID for the Oracle Bare Metal Cloud account.

--compartment_ocid
The compartment OCID within the tenancy of the Oracle Bare Metal Cloud Account.

Details

Modifies information about a cloud user.

racli alter credential_wallet

Changes the password on a wallet object in the Recovery Appliance.

Syntax

```
racli alter credential_wallet --wallet_path=PATH_TO_WALLET
```

Options

--wallet_path
Path the credential wallet will reside.

Details

Changes the password on a wallet object in the Recovery Appliance.

If no wallet_path is specified the default path is used /raacfs/raadmin/ra_wallet/.

Examples

This example alters a credential wallet in the defualt path

```
racli add credential_wallet
```

When prompted, enter the password for the new credential wallet.

```
Enter Wallet Password:
Re Enter Password:
```

racli alter keystore

Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.

Syntax

```
racli alter keystore { --initalize_key | --restart_db | --password | --
rekey }
```

Options

`--rekey`
Generates new TDE master encryption key (rekey) for the Recovery Appliance. Keys for protected databases are created using `dmbs_ra.key_rekey`.

`--password`
Specifies a new password for the keystore.

`--initialize_key`
Generates the initial TDE master key for the keystore.

`--restart_db`
Restarts the database stack after changes to the keystore to ensure they are honored at service restart.

Details

Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.

Examples

The following example restarts the database stack following a keystore change.

```
racli alter keystore --restart_db
```

racli alter parameter

Modifies the specified parameter whether it should be included or ignored as a check.

Syntax

```
racli alter parameter --type=db --parameter_name=NAME --ignore_check={ false |
true } [--force]
```

Options

`--type`
The type of the init parameter to be checked. Presently only "db" is supported.

`--parameter_name`
The name of the init parameter to be checked.

`--ignore_check`
Specifies whether to ignore a parameter (true) or to include the parameter (false) in a check of the database.

`--force`
Certain parameters are deemed important for database checks and are managed in an exclusion list to always be checked. The `--force` option bypasses the exclusion list so that the parameter can be ignored.

Details

Modifies the specified parameter whether it should be included or ignored as a check.

racli alter passphrase

.

Syntax

```
racli alter passphrase
```

Options

Modifies password used to access the Recovery Appliance.

Details:

- The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.

Details

Modifies password used to access the Recovery Appliance.

- The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.

racli alter rasys_user

Modifies the Recovery Appliance rasys user account password and management of the password.

Syntax

```
racli alter rasys_user { --init | --password }
```

Options

`--init`
Establishes management of the *rasys* user. This will be run for you during patch or upgrade in most cases.
In the event the *rasys* wallet goes away, this can be used to re establish rasys password management.

`--password`
Set a new *rasys* user password

Details

Modifies the Recovery Appliance rasys user account password and management of the password.

Examples

The following example changes the password for the `rasys` user account.

```
racli alter rasys_user --password
```

When prompted as shown, enter the new password.

```
Enter rasys Password:
Confirm new rasys Password:
```

The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.

racli alter vpc_user

Modifies an existing Recovery Appliance user account and changes its password.

Syntax

```
racli alter vpc_user --user_name=myVpcUserName
```

Options

`--user_name`
The Virtual Private Catalog user account for the Recovery Appliance

Details

Modifies an existing Recovery Appliance user account and changes its password.

The `user_name` specified in the command must be an existing Recovery Appliance user account.

The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used.

Examples

The following example changes the password for the `ravpc1` user account.

```
racli alter vpc_user --user_name=ravpc1
```

When prompted as shown, enter the new password.

```
ravpc1 New Password:
```

racli configure

Configure an object from operations performed by the Recovery Appliance.

Options

Details

Configure an object from operations performed by the Recovery Appliance.

racli configure archive_group

Configures archive log grouping.

Syntax

```
racli configure archive_group
```

Options

Details

Configures archive log grouping.

racli disable

Disables an object from operations performed by the Recovery Appliance.

Options

Details

Disables an object from operations performed by the Recovery Appliance. Object remains known to the Recovery Appliance for a quick enable later.

racli disable aide

Disables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.

Syntax

```
racli disable aide
```

Options

Details

Disables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.

AIDE is an Exadata utility that creates a database of files on the system and tracks changes. AIDE ensures file integrity and detects system intrusions. AIDE needs to be disabled during system installation and upgrades, and then afterward enabled again.

racli disable check

Disables a specific check object from check run operations performed on the Recovery Appliance.

Syntax

```
racli disable check { --check_name=CHECK_NAME }
```

Options

`--check_name`
Indicates the name of the check to be disabled.

Details

Disables a specific check object from check run operations performed on the Recovery Appliance.

racli disable cloud_location

Disables a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli disable cloud_location [ --location_name=CLOUD_LOCATION_NAME ]
```

Options

`--location_name`
Specifies the cloud location name to be disabled.

Details

Disables a cloud storage object for operations performed by the Recovery Appliance. The enable action defaults to **all** for cloud_location unless specified. This translates to 'pause' for `sbt_library` in a ready state.

racli disable keystore

Disables the keystore.

Syntax

```
racli disable keystore
```

Options

Details

Disables the keystore.

racli disable nfs

.

Syntax

```
racli disable nfs --name=myMountName
```

Options

Disables the specified Network File System (NFS).

Details

`--name`
Name of the NFS mount object.

racli enable

.

Options

Enables an object for operations performed by the Recovery Appliance.

Details

Enables an object for operations performed by the Recovery Appliance.

racli enable acfs

Enables the specified Automatic Storage Management Cluster (ACFS).

Syntax

```
racli enable acfs --name=myMountName
```

Options

`--name`
Name of the ACFS mount object.

Details

Enables the specified Automatic Storage Management Cluster (ACFS).

racli enable aide

Enables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.

Syntax

```
racli enable aide
```

Options

Details

Enables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.

AIDE is an Exadata utility that creates a database of files on the system and tracks changes. AIDE ensures file integrity and detects system intrusions. AIDE needs to be disabled during system installation and upgrades, and then afterward enabled again.

racli enable check

Enables a check on the Recovery Appliance.

Syntax

```
racli enable check { --check_name=CHECK_NAME };
```

Options

`--check_name`
Indicates the name of the check to be enabled on the Recovery Appliance.

Details

Enables a check on the Recovery Appliance.

racli enable cloud_location

Enables a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli enable cloud_location [ --location_name=CLOUD_LOCATION_NAME ]
```

Options

`--location_name`
Specifies the cloud location name to be enabled.

Details

Enables a cloud storage object for operations performed by the Recovery Appliance. The enable action defaults to **all** for cloud_location unless specified. This translates to 'resume' for `sbt_library` in a paused state.

racli enable keystore

Opens the keystore so that it can be used.

Syntax

```
racli enable keystore
```

Options

Details

Opens the keystore so that it can be used. This is the OKV endpoint keystore, also known as the OKV shared wallet.

racli enable nfs

Enables the specified Network File System (NFS).

Syntax

```
racli enable nfs --name=myMountName
```

Options

`--name`
Name of the NFS mount object.

Details

Enables the specified Network File System (NFS).

racli export

Used for exporting data.

Options

Details

Used for exporting data.

racli export appliance

Creates an export bundle for the Recovery Appliance.

Syntax

```
racli export appliance
```

Options

Details

Creates an export bundle for the Recovery Appliance.

racli install

Installs software.

Options

Details

Installs software.

racli install appliance

Installs software to the Recovery Appliance software.

Syntax

```
racli install appliance --step={ STEP_NUMBER | ALL } [--
import=IMPORT_BUNDLE_PATH]
```

Options

`--import`
Specifies the location of the RA export bundle to use for the installation.

`--step`
Indicate which step number of install appliance to run. Value range is 1-7 or "all" for all steps.
Step Numbers:

1. Validates and configures the hardware and platform.
2. Configures the operating system, including setting up file system properties, and users and groups.
3. Configures the `oracle` operating system user.
4. Creates the database file system (DBFS) mounts, configures TNS aliases, and sets up the required wallets for communication.
5. If a tape backup system is detected, validates and configures the appliance using details that were specified on the Tape Library Connectivity Page in the Oracle Exadata Deployment Assistant.
6. Configures backup of the Recovery Appliance catalog database and creates credentials for performing the backups. The catalog is backed up to a Recovery Appliance disk on an automated schedule.
7. Starts all Recovery Appliance services installed in the previous steps.

Details

Installs software to the Recovery Appliance software.

Before you run `racli install appliance`:

- Ensure that all current installation patches are stated in the `/u01/ra_install/linux-x64/WorkDir` directory.
- Ensure that the `ra_preinstall.pl` and `install.sh` scripts were run and successfully completed.

If a step fails, contact Oracle Support Services.

Examples

The following example runs only step 2 of the software installation.

```
racli install appliance --step=2
```

The following example runs all of the steps necessary to install the software based on the system and configuration settings.

```
racli install appliance
```

racli install okv_endpoint

Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance.

Syntax

```
racli install okv_endpoint
```

Options

Details

Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance. This is run after the a keystore has been activiated and is only run once. It adds all OKV endpoints associated with the Recovery Appliance; it applies to all nodes.

Examples

This example installs the okv_endpoints associated with each compute node of the Recovery Appliance.

```
racli install okv_endpoint


Enter New Keystore Password:
Confirm New Keystore Password:
Enter RA Wallet Password:
Wed December 23 20:14:40 2018:   Start: Install OKV End Point [compute_node1]
Wed December 23 20:14:43 2018:   End: Install OKV End Point [compute_node1]
Wed December 23 20:14:43 2018:   Start: Install OKV End Point [compute_node2]
Wed December 23 20:14:45 2018:   End: Install OKV End Point [compute_node2]
```

racli list

Lists information about objects given in the Recovery Appliance schema.

Options

Details

Lists information about objects given in the Recovery Appliance schema.

racli list check

List checks available on the Recovery Appliance.

Syntax

```
racli list check { --status=STATUS | --all } {--verbose}
```

Options

`--status`
List checks by STATUS. Valid STATUS are `enabled | disabled`.

`--all`
List names of all checks.

`--verbose`
List details of specified checks.

Details

List checks available on the Recovery Appliance.

racli list cloud_key

List details in short (default) or long format about a specified cloud key or all cloud keys.

Syntax

```
racli list cloud_key { --key_name=KEY_NAME | --all } [--long] [--fingerprint]
```

Options

`--key_name`
Specifies a key name associated with the cloud storage to list details.

`--fingerprint`
Displays the fingerprint associated with the cloud key.

`--all`
Lists details about all cloud keys known to the Recovery Appliance.

`--long`
Lists extended details about one or all cloud keys.

Details

List details in short (default) or long format about a specified cloud key or all cloud keys.

racli list cloud_location

Lists details about the cloud storage known to the Recovery Appliance.

Syntax

```
racli list cloud_location { --location_name=CLOUD_LOCATION_NAME | --all } [--
long]
```

Options

`--location_name`
Specifies a cloud location name to list details.

`--all`
Lists details about all cloud storage known to the Recovery Appliance.

`--long`
Lists extended details about one or all containers.

Details

Lists details about the cloud storage known to the Recovery Appliance.

Examples

The following example lists all known cloud_location objects.

```
racli list cloud_location --all
```

racli list cloud_user

List details about the specified cloud user, or all cloud users.

Syntax

```
racli list cloud_user { --user_name=USER_NAME | --all } [--long]
```

Options

`--user_name`
The name user wants associated with this particular cloud user.

`--all`
Lists details about all cloud users known to the Recovery Appliance.

`--long`
Lists extended details about one or all cloud users.

Details

List details about the specified cloud user, or all cloud users.

racli list compute_nodes

Lists the names of the compute nodes available on the Recovery Appliance.

Syntax

```
racli list compute_nodes
```

Options

Details

Lists the names of the compute nodes available on the Recovery Appliance.

racli list grid_home

Lists details about grid_home.

Syntax

```
racli list grid_home { --active | --inactive | --removed | --all }
```

Options

`--active`
List details about active grid_home.

`--inactive`
List details about all inactive grid_home.

`--removed`
List details about removed grid_home.

`--all`
List details about all grid_home.

Details

Lists details about grid_home.

racli list host

Lists details about all hosts available to the Recovery Appliance.

Syntax

```
racli list host
```

Options

Details

Lists details about all hosts available to the Recovery Appliance.

racli list nfs

Lists details about one or more Network File System mounts.

Syntax

```
racli list nfs { --mount_name=myName | --all }
```

Options

`--mount_name`
Name of NFS mount to be listed.

`--all`
Specifies that all NFS mounts should be listed.

Details

Lists details about one or more Network File System mounts.

racli list okv_endpoint

Lists details about the OKV end points known to the Recovery Appliance.

Syntax

```
racli list okv_endpoint [--all] [--long]
```

Options

`--all`
Lists all of the okv_endpoint(s) known to the Recovery Appliance.

`--long`
Lists extended details about the OKV end points known to the Recovery Appliance.

Details

Lists details about the OKV end points known to the Recovery Appliance.

Examples

The following example lists all of the OKV end points with extended details.

```
racli list okv_endpoint --all --long
```

racli list oracle_home

Lists details about oracle_home.

Syntax

```
racli list oracle_home { --active | --inactive | --removed | --all }
```

Options

`--active`
List details about active oracle_home.

`--inactive`
List details about all inactive oracle_home.

`--removed`
List details about removed oracle_home.

`--all`
List details about all oracle_home.

Details

Lists details about oracle_home.

racli list parameter

Lists the names and value of the init parameters for the Recovery Appliance.

Syntax

```
racli list parameter --type=db [--invalid | --valid | --all] [--long]
```

Options

`--type`
The type of the init parameter to be listed.

`--invalid`
Show only invalid init parameter names.

`--valid`
Show only valid init parameter names.

`--all`
Show all init parameter names.

`--long`
Show init parameter value as well as names.

Details

Lists the names and value of the init parameters for the Recovery Appliance.

racli list storage_cells

Lists names of storage cells available to the Recovery Appliance.

Syntax

```
racli list storage_cells
```

Options

Details

Lists names of storage cells available to the Recovery Appliance.

racli list tag

List diagnostic tags available on the Recovery Appliance.

Syntax

```
racli list tag { --tag=<TAG_NAME> | --all }
```

Options

`--tag`

`List diagnostics performed under the specific tag.`

`--all`

`List diagnostics performed by all tags.`

Details

List diagnostic tags available on the Recovery Appliance.

racli patch

The action to patch a specified object.

Options

Details

The action to patch a specified object.

racli patch appliance

Patch the Recovery Appliance Software.

Syntax

```
racli patch appliance --step=STEP_NUMBER [--ignore_incidents]
```

Options

`--step`
Indicate which step number of upgrade appliance to run. Value range is 1-2.
Step Numbers:

1. Runs the prechecks prior to upgrade. No downtime required. Does not change the software.
2. Performs the Patch. A full outage is required and is followed by a postcheck.

`--ignore_incidents`
Ignore active incidents (not recommended).

Details

Patch the Recovery Appliance Software.

racli remove

Removes objects from the Recovery Appliance schema.

Options

Details

Removes objects from the Recovery Appliance schema. Removed objects become unknown to the Recovery Appliance and cannot be quickly enabled later.

racli remove acfs

Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.

Syntax

```
racli remove acfs --mount_name=myName
```

Options

`--mount_name`
Name of ACFS mount to be removed.

Details

Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.

racli remove cloud_key

.

Syntax

```
racli remove cloud_key --key_name=KEY_NAME
```

Options

Removes the specified cloud key, but only if no cloud users are associated with the key.

Details

`--key_name`
The name of the public and private key to remove.

racli remove cloud_location

Removes the specified cloud storage from the Recovery Appliance, unless the cloud archive job template is tied to the library.

Syntax

```
racli remove cloud_location --location_name=LOCATION_NAME [--force]
```

Options

`--location_name`
Specifies the name of the cloud location to be removed.

`--force`
If a cloud storage job template is tied to the library, this option forces its removal anyway.

Details

Removes the specified cloud storage from the Recovery Appliance, unless the cloud archive job template is tied to the library.

racli remove cloud_user

Removes a cloud user.

Options

`--user_name`
The cloud user object name associated with this particular user.

Details

Removes a cloud user. A cloud user associated with a cloud location cannot be removed.

racli remove grid_home

Removes the specified grid home.

Syntax

```
racli remove grid_home --grid_home=inactiveGridHome
```

Options

`--grid_home`
Path of grid_home to be removed.

Details

Removes the specified grid home.

racli remove host

Removes the specified host.

Syntax

```
racli remove host --host_name=HOSTNAME
```

Options

`--host_name`
Indicates the name of the host to be removed.

Details

Removes the specified host.

racli remove keystore

Removes the specified keystore from the Recovery Appliance schema.

Syntax

```
racli remove keystore
```

Options

Details

Removes the specified keystore from the Recovery Appliance schema.

racli remove nfs

Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.

Syntax

```
racli remove nfs --mount_name=myName
```

Options

`--mount_name`
Name of NFS mount to be removed.

Details

Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.

racli remove oracle_home

Removes the specified oracle home.

Syntax

```
racli remove oracle_home --oracle_home=inactiveOracleHome
```

Options

`--oracle_home`
Path of oracle_home to be removed.

Details

Removes the specified oracle home.

racli remove osb_pieces

Removes all pieces that are no longer in the ZDLRA catalog from the OSB catalog.

Syntax

```
racli remove osb_pieces
```

Options

Details

Removes all pieces that are no longer in the ZDLRA catalog from the OSB catalog.

racli remove tag

Removes the specified diagnostic TAG from all nodes of the Recovery Appliance.

Syntax

```
racli remove tag --tag=TagName
```

Options

`--tag`
Name of the diagnostic tag to be removed.

Details

Removes the specified diagnostic TAG from all nodes of the Recovery Appliance.

racli remove vpc_user

Removes the specified user account from the Recovery Appliance metadata schema.

Syntax

```
racli remove vpc_user --user_name=myVPCUserName
```

Options

`--user_name`
The Virtual Private Catalog user for the Recovery Appliance.

Details

Removes the specified user account from the Recovery Appliance metadata schema.

- The specified user account must exist in the Recovery Appliance schema.
- The specified user account must not be associated with a protected database.

Examples

The following example removes the user account `ravpc1` from the metadata repository.

```
racli remove vpc_user --user_name=ravpc1
```

racli run

Performs an action for an object to the Recovery Appliance.

Options

Details

Performs an action for an object to the Recovery Appliance.

racli run check

Runs the enabled checks for the Recovery Appliance.

Syntax

```
racli run check { --check_name=CHECK_NAME | --group_name=GROUP_NAME | --all };
```

Options

`--all`
Indicates that all checks are run.

`--check_name`
Indicates the name of the check to be run.

--group_name
Indicates the name of the check group to be run.

Details

Runs the enabled checks for the Recovery Appliance.

racli run diagnostics

Runs and collects the required diagnostics on the Recovery Appliance.

Syntax

```
racli run diagnostics { --tag=TAG | --level=LEVEL | --path };
```

Options

--tag
Specifies the name of the tag of diagnostics to be collected, where default tags are:
`sar` 'system activity report'
`txt_dump` 'RA dump txt file'
`dp_dump` 'RA dump datapump'
`ord_wait` 'ORDERING_WAIT analysis'
`corruption` 'corruption'
`racli_logs` 'gather racli logs from each node'
`ra_version` 'get RA version information from each node'
`str_cfg` 'get store and config files from each node'
`history` 'get root and oracle history from each node'
`network` 'get network information from each node'
`obtooldiag` 'run obtool gathering useful information'
`osblogs` 'get osb logs from each node'
`osbxcr` 'get /usr/etc/ob/xcr from each node'
`osbstate` 'get backup/admin/state from each node'
`alerts` 'get ASM and RDBMS alert.logs from each node'
`awr` 'generate awr dump for last 7 days'
And collection tags:
ra_low:

```
  [1] Generates a System Activity Report on the Recovery Appliance.

  [2] Gets the version information from each node of the Recovery
Appliance.
```

ra_med:

```
  [1] Generates a System Activity Report on the Recovery Appliance.

  [2] Performs a text format dump of the Recovery Appliance.

  [3] Gets the version information from each node of the Recovery
Appliance.
```

ra_high:

```
  [1] Generates a System Activity Report on the Recovery Appliance.

  [2] Performs a datapump dump of the Recovery Appliance.

  [3] Gets the version information from each node of the Recovery Appliance.
```

osb_low:

```
  [1] Gets OSB information through obtool on the Recovery Appliance.

  [2] Gathers OSB logs from all nodes of the Recovery Appliance.
```

osb_med:

```
  [1] Gets OSB information through obtool on the Recovery Appliance.

  [2] Gathers OSB logs from all nodes of the Recovery Appliance.

  [3] Gathers OSB xcr files from all nodes of the Recovery Appliance.
```

osb_high:

```
  [1] Gets OSB information through obtool on the Recovery Appliance.

  [2] Gathers OSB logs from all nodes of the Recovery Appliance.

  [3] Gathers OSB xcr files from all nodes of the Recovery Appliance.

  [4] Gathers OSB state files from the current node of the Recovery Appliance.
```

`--level`
Alternatively, if a tag doesn't exist the level can be specified directly, where:
level 1 System activity report.
level 2 RA dump txt file.
level 4 RA dump datapump.
level 8 ORDERING_WAIT analysis.
level 16 Corruption.
level 32 Gather racli logs from each node.
level 64 Get RA version information from each node.
level 128 Get store and config files from each node.
level 256 Get root and oracle history from each node.
level 512 Get network information from each node.
level 1024 Run obtool gathering useful information.
level 2048 Get osb logs from each node.
level 4096 Get /usr/etc/ob/xcr from each node.
level 8192 Get backup/admin/state from each node.
level 16384 Get all nodes ASM and ZDLRA alert.logs.
level 32768 Generate awr dump for last 7 days.
The levels are additive, so to gather a system activity report and RA dump txt file you would use --level=3

`--path`
Specfies the location on the Recovery Appliance where the diagnostic bundle will be created, default being radump.

Details

Runs and collects the required diagnostics on the Recovery Appliance.

Examples

The following example collects racli logs from all nodes and creates a bundle in default location /radump

```
racli run diagnostics --tag=racli_logs
```

The following example using the level directly, runs and bundles system activity report and alert.logs and creates a bundle in path /tmp.

```
racli run diagnostics --level=16385 --path=/tmp
```

racli start

The action to start a specified service (object).

Options

Details

The action to start a specified service (object).

racli start nfs

Start NFS services of the Recovery Appliance.

Syntax

```
racli start nfs --name=myMountName
```

Options

`--name`
Name of the NFS mount to start.

Details

Start NFS services of the Recovery Appliance.

racli status

Check status of objects associated with the Recovery Appliance.

Options

Details

Check status of objects associated with the Recovery Appliance.

racli status aide

Returns the status of the Advanced Intrusion Detection Environment (AIDE).

Syntax

```
racli status aide
```

Options

Details

Returns the status of the Advanced Intrusion Detection Environment (AIDE).

racli status appliance

Returns a summary of the status of the appliance.

Syntax

```
racli status appliance
```

Options

Details

Returns a summary of the status of the appliance.

racli status cloud_location

Lists the status of a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli status cloud_location [ --location_name=CLOUD_LOCATION_NAME ] [--long]
```

Options

`--location_name`
Specifies the name of the cloud location for status.

`--long`
Lists extended details about cloud storage known to the Recovery Appliance.

Details

Lists the status of a cloud storage object for operations performed by the Recovery Appliance. This action defaults to **all** for cloud_location unless specified.

racli status disk_group

Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.

Syntax

```
racli status disk_group
```

Options

Details

Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.

Run this command before running the `racli update storage_location` command to ensure that the ASM rebalance completed.

All rebalance must complete before Griddisks and Storage Spaces are added to a Recovery Appliance.

racli status keystore

Returns the status of the keystore in use by the Recovery Appliance.

Syntax

```
racli status keystore
```

Options

Details

Returns the status of the keystore in use by the Recovery Appliance.

racli status nfs

Returns the status of the Network File Service (NFS).

Syntax

```
racli status nfs --name=myMountName
```

Options

Details

Returns the status of the Network File Service (NFS).

racli status okv_endpoint

Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.

Syntax

```
racli status okv_endpoint
```

Options

Details

Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.

racli stop

The action to stop a specified service (object).

Options

Details

The action to stop a specified service (object).

racli stop nfs

Stop NFS services of the Recovery Appliance.

Syntax

```
racli stop nfs --name=myMountName
```

Options

`--name`
Name of the ACFS mount to start.

Details

Stop NFS services of the Recovery Appliance.

racli uninstall

Backs out a portion of the installation.

Options

Details

Backs out a portion of the installation.

racli uninstall appliance

Backout a step during install.

Syntax

```
racli uninstall appliance --step={ STEP_NUMBER | ALL }
```

Options

`--step`
Indicate which step number of uninstall appliance to run. Value range is 1-7 or "all" for all steps.
Step Numbers:

1. Unconfigures the hardware and platform.
2. Unconfigures the operating system, including setting up file system properties, and users and groups.
3. Unconfigures the Oracle operating system user.
4. Removes file system mounts, unconfigures TNS aliases, and invalidates wallets for communication.
5. If a tape backup system is detected, unconfigures the resources used for tape backup.
6. Unconfigures backup of the Recovery Appliance catalog database and removes credentials for performing the backups.
7. Stops all Recovery Appliance services installed.

Details

Backout a step during install.

Generally used if you found an error and have to redo a step after correcting. You must run the uninstallation in the reverse numerical order.

> **Note:**
>
> If a failed installation step is the reason that you plan to run the uninstall command, it is highly recommended that you contact Oracle Support Services.

Examples

The following example undoes only step 3 of the installation procedure.

```
racli uninstall appliance --step=3
```

The following undoes all of the steps of the installation procedure in reverse sequence.

```
racli uninstall appliance
```

racli uninstall okv_endpoint

Uninstalls the specified OKV end point.

Syntax

```
racli uninstall okv_endpoint [--force]
```

Options

`--force`
When specified, forces the uninstallation of the OKV end point.
Otherwise, a check is done for active storage locations utilizing the OKV end point. If any are in use, a warning is issued and uninstall fails.

Details

Uninstalls the specified OKV end point.

Examples

This example forces uninstallation of the OKV end points associated with the Recovery Appliance.

```
racli uninstall okv_endpoint --force


Thu August 17 06:03:55 2018: Start: Remove OKV End Point
Thu August 17 06:03:55 2018:
  WARNING:
    Found active cloud locations.
    Removing the OKV endpoint will mean this Recovery Appliance can't recover
      backups at cloud locations without an endpoint configured. Thu August 17
06:03:55 2018:   Start: Remove OKV End Point [node1]
Thu August 17 06:03:55 2018:   End: Remove OKV End Point [node1]
Thu August 17 06:03:55 2018:   Start: Remove OKV End Point [node2]
Thu August 17 06:03:55 2018:   End: Remove OKV End Point [node2]
Thu August 17 06:03:55 2018:   End: Remove OKV End Point
```

racli update

Update a component of the Recovery Appliance.

Options

Details

Update a component of the Recovery Appliance.

racli update aide

Updates the Advanced Intrusion Detection Environment (AIDE).

Syntax

```
racli update aide
```

Options

Details

Updates the Advanced Intrusion Detection Environment (AIDE).

AIDE is an Exadata utility that creates a database of files on the system and tracks changes. AIDE ensures file integrity and detects system intrusions.

racli update parameter

Update init parameters in the Recovery Appliance.

Syntax

```
racli update parameter --type=db [--restart_db]
```

Options

`--type`
The type of init parameter to be updated. Presently only "db" is supported.

Details

Update init parameters in the Recovery Appliance.

racli update storage_location

Update RA pool data file with the most current number of storage cells.

Syntax

```
racli update storage_location --cells=cell1,cell2,...celln
```

Options

Details

Update RA pool data file with the most current number of storage cells.

racli update support_tools

Updates the AHF and Exachk versions using the staged binaries zip files.

Syntax

```
racli update support_tools {--ahf_file=<AHF_FILE_NAME>} {--
exachk_file=<EXACHK_FILE_NAME>}
```

Options

`--ahf_file`
Updates AHF version the the provided AHF version.

`--exachk_file`
Updates Exachk version the the provided Exachk version.

Details

Updates the TFA and Exachk versions using the staged binaries zip files.

racli upgrade

The action to upgrade a specified object to the next version.

Options

Details

The action to upgrade a specified object to the next version.

racli upgrade appliance

Upgrade the Recovery Appliance Software Versions.

Syntax

```
racli upgrade appliance --step=STEP_NUMBER [--ignore_incidents]
```

Options

`--step`
Indicate which step number of upgrade appliance to run. Value range is 1-5.
Step Numbers:

1. Runs the prechecks prior to upgrade. No downtime required. Does not change the software.
2. Stages the new software stack. No downtime required. Does not change the software.
3. Performs the upgrade using the staged items. A full outage is required; it is followed by a postcheck and cleanup of the old files.

4. Migrate DBFS to ACFS. All files residing on DBFS will be migrated to new ACFS.
5. Upgrade OSB from 12.1.x.x.x to 12.2.0.1.0..

`--ignore_incidents`
Ignore active incidents (not recommended).

Details

Upgrade the Recovery Appliance Software Versions.

racli version

Provides version information for software running on the Recovery Appliance.

Syntax

```
racli version
```

Options

rarpm version
The ra_automation rpm currently installed on the Recovery Appliance.

rdbms version
The database version currenly installed on the Recovery Appliance.

transaction
The specific transaction associated with the version running on the Recovery Appliance.

zdlra version
The base software version installed and running on the Recovery Appliance.

Details

Provides version information for software running on the Recovery Appliance.

The versioning information provided includes:

B
Error Messages

This appendix contains the hardware messages for Recovery Appliance. They are identical to the messages for Oracle Exadata Database Machine.

This appendix contains the following sections:

> **See Also:**
>
> *Oracle Database Error Messages* for information about Oracle Database SNMP error messages

Contacting Oracle Support Services

Some messages recommend contacting Oracle Support Services to report a problem. To report a problem, sign in to My Oracle Support at `http://support.oracle.com` and log a service request. You might also want to contact Oracle Support when Auto Service Request submits a service request.

When you contact Oracle Support Services, have the following information available:

- The hardware, operating system, and release number of the operating system running Oracle Database.

- The complete release number of Oracle Database, such as release 12.1.0.2.
- All Oracle programs (with release numbers) in use when the error occurred.
- If you encountered one or more error codes or messages, then the exact code numbers and message text, in the order in which they appeared.
- The problem severity, according to the following codes:
  - 1: Program not usable. Critical effect on operations.
  - 2: Program usable. Operations severely restricted.
  - 3: Program usable with limited functions. Not critical to overall operations.
  - 4: Problem circumvented by customer. Minimal effect, if any, on operations.

You will also be expected to provide the following:

- The name of your organization
- Your name
- Your Oracle Support ID number
- Your telephone number
- Rack master serial number

See "#unique_443" for links to My Oracle Support notes that are related to Recovery Appliance.

Generic Alert Messages: HALRT-00001 to HALRT-00500

HALRT-00002, Chassis Intrusion - Physical Security Violation
Cause: Detected a Physical Security Violation on the chassis.

Action: See "HALRT-10001".

HALRT-00003, Chassis Intrusion (Physical Security Violation) Event Cleared
Cause: Detected a clear event of Physical Security Violation on the chassis.

Action: See "HALRT-10001".

HALRT-00004, CPU Internal Error
Cause: Detected an Internal CPU error.

Action: See "HALRT-10025".

HALRT-00005, CPU Internal Error Cleared
Cause: Detected a clear event for CPU internal error.

Action: See "HALRT-10001".

HALRT-00006, CPU Configuration Error
Cause: Detected a CPU configuration error.

Action: See "HALRT-10025".

HALRT-00007, CPU Configuration Error Cleared
Cause: Detected a clear event for a CPU configuration error.

Action: See "HALRT-10001".

HALRT-00008, CPU Presence (Processor Presence detected)
Cause: Detected presence event for a new CPU.

Action: See "HALRT-10025".

HALRT-00009, CPU Not Present (Processor Not Present)
Cause: Detected a CPU is not present.

Action: See "HALRT-10001".

HALRT-00010, System Event Log (SEL) Full (Logging Disabled)
Cause: System event log detected to be full.

Action: See "HALRT-10001".

HALRT-00011, ASR (Automatic System Recovery) Timer Expired
Cause: Automatic System Recovery timer expired.

Action: See "HALRT-10001".

HALRT-00012, ASR (Automatic System Recovery) Reset Occurred
Cause: Automatic System Recovery reset occurred.

Action: See "HALRT-10001".

HALRT-00013, ASR (Automatic System Recovery) Power Down Occurred
Cause: Automatic System Recovery was powered down.

Action: See "HALRT-10001".

HALRT-00014, ASR (Automatic System Recovery) Power Cycle Occurred
Cause: Automatic System Recovery power was cycled.

Action: See "HALRT-10001".

HALRT-00015, {0} Sensor Threshold state change: Sensor Component Name: {1} Sensor State: {2} Sensor Threshold Type: {3} Sensor Threshold Severity: {4} Sensor Value: {5} Sensor Threshold Value: {6} Sensor Additional Info: {7}
Cause: The sensor threshold state has changed.

Action: See "HALRT-10029, A generic sensor has detected an error.".

HALRT-00016, A {0} component is suspected of causing a fault. Component Name: {1} Fault class: {2} Fault certainty: {3} Fault message: {4} Trap Additional Info: {5}
Level Critical

Cause: A component might have caused a fault to occur.

Action: See "HALRT-10030".

HALRT-00017, A {0} component fault has been cleared. Component Name: {1} Trap Additional Info: {2}
Level Clear

Cause: The component fault has cleared.

Action: See "HALRT-10001".

HALRT-00018, A {0} sensor has detected an error. Component Name: {1} Trap Additional Info: {2}
Cause: An error was detected by a sensor.

Action: See "HALRT-10029, A generic sensor has detected an error.".

HALRT-00019, A {0} sensor has returned to its normal state. Component Name: {1} Trap Additional Info: {2}
Cause: The sensor has returned to its normal state.

Action: See "HALRT-10001".

HALRT-00020, A field replaceable unit at {0} has been inserted into the system.
Cause: A unit was inserted into the system.

Action: See "HALRT-10001".

HALRT-00021, A field replaceable unit at {0} has been removed from the system.
Cause: A unit was removed from the system.

Action: See "HALRT-10001".

HALRT-00022, The component operational state has changed or is changing to disabled. Component Name: {0} Disable Reason: {1} Trap Additional Info: {2}
Cause: The operational state for a component has changed.

Action: See "HALRT-10001".

HALRT-00023, The component operational state has changed or is changing to enabled. Component Name: {0} Trap Additional Info: {1}
Cause: The operational state for a component has changed.

Action: See "HALRT-10001".

HALRT-00024, An error has been detected in the period before the operating system has taken control of the system.
Cause: An error was detected.

Action: See "HALRT-10032".

HALRT-00025, The high availability component identified by {0} has changed state. Old HA State: {1} New HA State: {2} Trap Additional Info: {3}
Cause: The state for a high availability component has changed.

Action: See "HALRT-10030".

HALRT-00026, An intrusion sensor has detected that someone may have physically tampered with the system.
Cause: A sensor has detected a change to the system.

Action: See "HALRT-10001".

HALRT-00027, The hard drive has changed state. Component Name: {0} New State: {1}
Cause: The state of the hard drive has changed.

Action: See "HALRT-10001".

HALRT-00028, This is a test trap.
Cause: User generated a test trap, and an alert was created.

Action: See "HALRT-10001".

HALRT-00030, A {0} component fault has been cleared.
Cause: A component fault was cleared.

Action: This is an informational message, and no action is required.

Temperature Alert Messages: HALRT-00513 to HALRT-00600

HALRT-00513, Under-Temperature Warning (Lower non-critical, going low).
Cause: The temperature was below the specified range.

Action: See "HALRT-10015"

HALRT-00514, Under-Temperature Warning Cleared.
Cause: The low temperature warning was cleared.

Action: See "HALRT-10001".

HALRT-00515, Under-Temperature Problem (Lower Critical - going low)
Cause: A low temperature problem was detected.

Action: See "HALRT-10004".

HALRT-00516, Under-Temperature Problem Cleared
Cause: The low temperature problem was cleared.

Action: See "HALRT-10001".

HALRT-00517, Over-Temperature warning (Upper non-critical, going high)
Cause: The temperature was above the specified range.

Action: "HALRT-10015".

HALRT-00518, Over-Temperature warning Cleared
Cause: The high temperature warning was cleared.

Action: See "HALRT-10001".

HALRT-00519, Over-Temperature Problem (Upper Critical - going high)
Cause: A high temperature was detected.

Action: "HALRT-10004".

HALRT-00520, Over-Temperature Problem Cleared
Cause: The high temperature problem was cleared.

Action: See "HALRT-10001".

HALRT-00521, CPU Thermal Trip (Over Temperature Shutdown)
Cause: The CPU high temperature shut down was tripped.

Action: See "HALRT-10004".

HALRT-00522, CPU Thermal Trip (Over Temperature Shutdown) Cleared
Cause: The CPU high temperature shut down was cleared.

Action: See "HALRT-10001".

Fan Alert Messages: HALRT-00616 to HALRT-00700

HALRT-00616, Generic Critical Fan Failure
Cause: A fan has failed.

Action: See "HALRT-10005".

HALRT-00617, Generic Critical Fan Failure Cleared
Cause: A fan failure was cleared.

Action: See "HALRT-10001".

HALRT-00618, Fan Redundancy Lost
Cause: A redundant fan has stopped functioning.

Action: See "HALRT-10001".

HALRT-00619, Fan redundancy Has Returned to Normal
Cause: A redundant fan has started to function normally.

Action: See "HALRT-10001".

Power Supply/Converters Alert Messages: HALRT-00717 to HALRT-00800

HALRT-00717, Power Supply Redundancy Lost
Cause: A redundant power supply has stopped functioning.

Action: See "HALRT-10001".

HALRT-00718, Power Supply Redundancy Has Returned to Normal
Cause: A redundant power supply has started to function normally.

Action: See "HALRT-10001".

HALRT-00719, Power Supply Inserted
Cause: A power supply was inserted into the system.

Action: See "HALRT-10001".

HALRT-00720, Power Supply Removed
Cause: A power supply was removed from the system.

Action: "HALRT-10007".

HALRT-00721, Power Supply Failure
Cause: A power supply has failed.

Action: "HALRT-10007".

HALRT-00722, Power Supply Failure Cleared
Cause: A power supply failure was cleared.

Action: See "HALRT-10001".

HALRT-00723, Power Supply Warning
Cause: A power supply problem has set off a warning.

Action: "HALRT-10007".

HALRT-00724, Power Supply Warning Cleared
Cause: A power supply warning was cleared.

Action: See "HALRT-10001".

HALRT-00725, Power Supply AC Lost
Cause: A power supply has lost AC power.

Action: Examine the power supply. Replace the power supply if necessary.

HALRT-00726, Power Supply AC Restored
Cause: A power supply has had it AC power restored.

Action: See "HALRT-10001".

Network Alert Messages: HALRT-00912 to HALRT-01000

HALRT-00912, Infiniband port {0} indicates invalid state: {1}
Cause: InfiniBand device indicated an invalid state.

Action: See "HALRT-11001".

HALRT-00913, InfiniBand port {0} indicates invalid rate: {1} Gps
Cause: InfiniBand device indicated invalid rate.

Action: See "HALRT-11003".

HALRT-00914, InfiniBand port status is OK
Cause: InfiniBand device status was all right.

Action: See "HALRT-10001".

HALRT-00915, InfiniBand port {0} is showing non-zero error counts
Cause: The error counts for `RcvErrors` and `SymbolErrors` on the specified port were non-zero. A non-zero error count can be caused by a bad cable, or a cable that is not properly seated, or by a bad port on the HCA or switch. Manually disabling and enabling InfiniBand ports or resetting InfiniBand switches can also lead to these errors.

Action: See "HALRT-11003".

Hard Disk and Flash Disk Alert Messages: HALRT-01015 to HALRT-01200

These alert messages are in the range HALRT-01015 through HALRT-01200.

HALRT-01015, Physical drive inserted on Adapter: {0} Device Id: {1}, Enclosure ID: {2}, Slot ID: {3}
Cause: The physical drive was inserted.

Action: See "HALRT-10001".

HALRT-01016, Physical drive removed on Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}
Cause: The physical drive was removed.

Action: See "HALRT-10040".

HALRT-01017, Physical drive state changed from {0} to {1}. Adapter: {2} Device ID: {3}, Enclosure ID: {4}, Slot ID: {5}
Cause: The physical drive status has changed.

Action: See "HALRT-10001".

HALRT-01018, FRU changed. Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}, Enclosure ID: {4} FRU {5}
Cause: The field replaceable unit (FRU) was changed on the adapter.

Action: See "HALRT-10001".

HALRT-01020, Physical drive not supported. Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}
Cause: The physical drive is not supported on the adapter.

Action: See "HALRT-10040".

HALRT-01021, Physical drive not certified. Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}
Cause: The physical drive was not certified.

Action: See "HALRT-10040".

HALRT-01024, Physical {0} disk error count is increasing. Disk: {1}, Serial Number: {2}, Slot Number: {3}
Level Critical

Cause: The disk error count increased during the periodic disk checks.

Action: Monitor the physical disk errors and replace the disk necessary.

HALRT-01025, Physical {0} disk error count is stable. Disk: {1}, Serial Number: {2}, Slot Number: {3}
Level Clear

Cause: The disk error count did not increase after the last disk check.

Action: This is an informational message, and no action is required.

HALRT-01027, Flash celldisk {0} write performance is degraded. The supercap module on the Flash card may have failed. Flash disk name : {1} Disk serial number: {2} Slot number: {3}
Cause: The write performance for the flash cell disk was degraded.

Action: Replace the flash disk, flash PCI card, or Energy Storage Module (ESM). If all disks on the PCI card show degraded performance, then check the ESM LED on the PCI card. It is visible from the back of the cell chassis. If this LED is amber, then replace the ESM. If the disks being replaced have persistent data, then rebalance the data out of these disks. Next, replace the flash disk. If persistent data was on the flash disks, then rebalance then data back into the disks. The flash cache portion provided by this flash disk will be disabled automatically.

HALRT-01028, Flash celldisk {0} write performance problem has been resolved
Cause: The performance problem for the flash cell disk was resolved.

Action: This is an informational message, and no action is required.

HALRT-01029, Physical disk {0} could not be set to ONLINE automatically. Physical disk {0} firmware could not be upgraded.
Level Critical

Cause: The physical disk firmware could not be upgraded.

Action: Upgrade the firmware on the physical disks. Restart MS to put the disk online.

HALRT-01032, Physical disk {0} firmware successfully upgraded
Cause: The physical disk firmware was upgraded.

Action: This is an informational message, and no action is required.

HALRT-01035, Physical disks {0} firmware could not be upgraded. Inserted physical disks will be accepted because MS_ACCEPT_DISKFW_DIFFERENT is set to true.
Level Informational

Cause: The firmware upgrade failed.

Action: This is an informational message.

HALRT-01036, Oracle Exadata discovered the following problems in the device utilities: {0}
Cause: A problem has been discovered. The problem is one of the following:

- Output has duplicate name.
- Output has duplicate serial number.
- Output has duplicate OS device name.
- Cannot find device name for the disk with target ID.

Action: The device monitoring utilities had unexpected outputs. A resolved alert will be sent when the problem is resolved. If the problem persists, then contact Oracle Support Services.

HALRT-01037, {0} output has duplicate name: {1}
Cause: The output had a duplicate name.

Action: See "HALRT-10073".

HALRT-01038, {0} output has duplicate serial number: {1}
Cause: The output had a duplicate serial number.

Action: See "HALRT-10073".

HALRT-01039, {0} output has duplicate OS device name: {1}
Cause: The output had a duplicate operating system device name.

Action: "HALRT-10073".

HALRT-01040, Cannot find device name for the disk with target ID: {0}
Cause: The device was not found.

Action: "HALRT-10073".

HALRT-01041, The problems in the device utilities have been resolved.
Cause: The problem was resolved.

Action: This is an informational message, and no action is required.

HALRT-01042, {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Warning

Cause: A disk was replaced.

Action: The grid disks and cell disk are re-created on the new physical disk. If the grid disks were part of an Oracle ASM disk group, then they will be added back to the disk group, and the data will be rebalanced on them, based on the disk group redundancy and the ASM_POWER_LIMIT parameter.

HALRT-01043, {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10} Flash Cache: {11} Flash Log {12}
Level Warning

Cause: The flash disk was replaced.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then flash cache will be reenabled automatically.
- If the flash disk in this slot was previously used for flash log, then flash log will be reenabled automatically.
- If the flash disk in this slot was previously used for grid disks, then Oracle ASM will re-create the grid disks and rebalance the data automatically.

HALRT-01044, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Critical

Cause: A system disk failed.

Action: Locate the affected drive and replace it as soon as possible. A white LED is lit on the affected storage server, and both an amber service action LED and a blue "OK to Remove" LED are lit on the drive. Oracle ASM rebalance will automatically restore the data redundancy.

HALRT-01045, {0} entered predictive failure status. {1} was replaced. Status: {2} Manufacturer : {3} Model Number : {4} Size: {5} Serial Number : {6} Firmware: {7} Slot Number : {8} Cell Disk: {9} Grid Disk: {10}
Level Critical

Cause: The system hard disk entered predictive failure status.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent and a blue OK to Remove LED is lit on the drive when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01046, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Critical

Cause: The hard disk failed.

Action: Locate the failed disk and replace it as soon as possible. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive.

HALRT-01047, {0} entered predictive failure status. {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Critical

Cause: The data hard disk entered predictive failure status.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent and a blue OK to Remove LED is lit on the drive when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01048, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Critical

Cause: The flash disk failed.

Action: Replace the flash disk. A white LED is lit on the affected server. If the flash disk is used for flash cache, then flash cache is disabled on this disk, thus reducing the effective flash cache size. If the flash disk is used for flash log, then flash log is disabled on this disk, thus reducing the effective flash log size. If the flash disk is used for grid disks, then Oracle ASM rebalance automatically restores the data redundancy.

HALRT-01049, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Critical

Cause: The system hard disk failed.

Action: Locate the failed drive, and replace it as soon as possible. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance automatically restores the data redundancy.

HALRT-01050, {0} entered predictive failure status. {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Cause: The system hard disk entered predictive failure.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01051, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}.
Level Critical

Cause: The data hard disk failed.

Action: Replace the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically restore the data redundancy.

HALRT-01052, {0} entered predictive failure status. {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}
Level Critical

Cause: The data hard disk entered predictive failure.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent and a blue "OK to Remove" LED is lit on the drive when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01053, The Oracle ASM drop operation on the following disk failed due to ORA-{0} ASM Disk Name: {1} ASM Diskgroup: {2} Grid Disk: {3} Cell Disk: {4} Physical Disk: {5} Physical Disk Status: {6}
Level Critical

Cause: The grid disk was not dropped from the Oracle ASM disk group, because of the specified reason.

Action: None. The storage server software continues to retry the command, and sends an alert when the command succeeds.

HALRT-01054, The Oracle ASM drop operation on the following disk failed due to ORA-{0} ASM Disk Name: {1} ASM Diskgroup: {2} Grid Disk: {3} Cell Disk: {4} Physical Disk: {5} Physical Disk Status: {6}
Level Critical

Cause: The Oracle ASM `DROP` command failed, because of the specified reason.

Action: Retry the `DROP` command manually on the Oracle ASM disk.

HALRT-01055, The Oracle ASM rebalance operation on the following disk failed due to ORA-{0} ASM Disk Name: {1} ASM Diskgroup: {2} Grid Disk: {3} Cell Disk: {4} Physical Disk: {5} Physical Disk Status: {6}
Level Critical

Cause: The Oracle ASM rebalance operation failed, because of the specified reason.

Action: This is an informational message, and no action is required.

HALRT-01056, The Oracle ASM rebalance operation on the following disk resumed ORA-{0} ASM Disk Name: {1} ASM Diskgroup: {2} Grid Disk: {3} Cell Disk: {4} Physical Disk: {5} Physical Disk Status: {6}
Level Warning

Cause: The Oracle ASM rebalance operation resumed.

Action: This is an informational message, and no action is required.

HALRT-01057, The Oracle ASM add operation on the following disk failed due to ORA-{0} ASM Disk Name: {1} ASM Diskgroup: {2} Grid Disk: {3} Cell Disk: {4} Physical Disk: {5} Physical Disk Status: {6}
Level Critical

Cause: The Oracle ASM operation to add the disk failed because of the specified reason.

Action: None. The storage server software continues to retry the command, and sends an alert when the command succeeds.

HALRT-01058, The Oracle ASM add operation on the following disk failed due to ORA-{0} ASM Disk Name: {1} ASM Diskgroup: {2} Grid Disk: {3} Cell Disk: {4} Physical Disk Name: {5} Physical Disk Status: {6}
Level Critical

Cause: The Oracle ASM operation to add the disk failed, because of the specified reason.

Action: Retry the command manually on the Oracle ASM disk.

HALRT-01059, Oracle Exadata completed the rebalance operation triggered by the add operation on the following ASM disks. ASM Disk Name: {0} ASM Diskgroup: {1} Grid Disk: {2} Cell Disk: {3} Physical Disk: {4} Physical Disk Status : {5}
Level Clear

Cause: The rebalance operation completed.

Action: This is an informational message, and no action is required.

HALRT-01060, Oracle Exadata completed the rebalance operation triggered by the drop operation on the following ASM disks. ASM Disk Name: {0} ASM Diskgroup: {1} Grid Disk: {2} Cell Disk: {3} Physical Disk: {4} Physical Disk Status : {5}
Level Warning

Cause: The rebalance operation to the Oracle ASM disk group completed.

Action: This is an informational message, and no action is required.

HALRT-01061, {0} is ready to use. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Clear

Cause: The flash disk was ready.

Action: This is an informational message, and no action is required.

HALRT-01062, {0} is ready to use. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Clear

Cause: The hard disk was ready.

Action: This is an informational message, and no action is required.

HALRT-01063, The Oracle ASM rebalance operation on the following disk failed as ASM rebalance operation is disabled. ASM Disk Name: {0} ASM Diskgroup: {1} Grid Disk: {2} Cell Disk: {3} Physical Disk: {4} Physical Disk Status : {5}
Level Critical

Cause: The rebalance operation failed, because the Oracle ASM rebalance operation is disabled.

Action: Enable the rebalance operation, and retry the command.

HALRT-01106, All Logical drives are in WriteThrough caching mode. HDD disk controller battery needs to be replaced. Please contact Oracle Support.
Cause: The logical drive was found.

Action: See "HALRT-10001".

HALRT-01107, Disk drive {0} is in WriteThrough caching mode. Write Performance on this drive will be affected. Please contact Oracle Support.
Cause: The disk drive was lost.

Action: See "HALRT-10054".

HALRT-01108, Logical drive status changed. LUN: {0}. Status: {1}. Physical {2} disk: {3}. Slot Num: {4}. Serial Number: {5}. Cell Disk: {6}. Grid Disks: {7}
Cause: The logical drive status has changed.

Action: See "HALRT-10001".

HALRT-01109, Disk controller battery temperature exceeded the threshold of {0} Celsius. All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. Batter ySerial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge {5} Ambient Temperature: {6}
Cause: The temperature exceeded the threshold.

Action: Verify that the ambient temperature is below 32 Celsius and the fans are working.

HALRT-01110, Logical drive found. It was empty. Failed creating cell disk and grid disks on it. LUN: {0}. Status: {1}. Physical: {2}. Slot Num: {3}. Serial Number: {4}
Cause: The logical drive was found.

Action: See "HALRT-10001".

HALRT-01115, BGI completed with uncorrectable errors on Adapter {0} VD Target {1}
Cause: BGI completed with errors.

Action: See "HALRT-10001".

HALRT-01116, BGI failed on Adapter {0} VD Target {1}
Cause: BGI failed on adapter.

Action: See "HALRT-10039".

HALRT-01117, LUN state changed on Adapter {0} VD Target {1} from {2} to {3}
Cause: The LUN state was changed.

Action: See "HALRT-10001".

HALRT-01118, LUN bad block table cleared on Adapter {0} VD Target {1}
Cause: The LUN bad block table was cleared.

Action: See "HALRT-10001".

HALRT-01123, {0} is rejected due to incorrect disk model. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Critical

Cause: The hard disk or LUN was removed. The storage server cannot use this hard disk because the disk model is not supported.

Action: Remove the hard disk and replace it with a supported disk.

> **See Also:**
>
> "Maintaining the Physical Disks of Storage Servers".

HALRT-01124, Hard Disk replaced. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8}
Level Warning

Cause: The hard disk was replaced.

Action: This is an informational message, and no action is required.

HALRT-01126, The full charge capacity of the disk controller battery has fallen below the minimum threshold of {0} mAh. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be somewhat reduced. The flash drives are not affected.Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: The full charge capacity of the battery was below the threshold.

Action: Replace the disk controller battery.

HALRT-01127, All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: The cause could not be determined.

Action: Contact Oracle Support Services.

HALRT-01128, Cell disk was missing and is recreated on the {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Cause: The cell disk was missing, and has been recreated on the disk.

Action: See HALRT-10001"HALRT-10001".

HALRT-01130, {0} status changed to poor performance. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Cause: The hard disk or flash disk status changed to poor performance.

Action: See "Maintaining the Physical Disks of Storage Servers".

HALRT-01131, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Warning

Cause: The hard disk or flash disk status changed to normal.

Action: This is an informational message, and no action is required.

HALRT-01132, {0} removed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Error Count: {10} Last Failure: {11}
Level Critical

Cause: The hard or flash disk was removed.

Action: This is an informational message, and no action is required.

HALRT-01133, {0} status changed to offline. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}.
Level Critical

Cause: The hard or flash disk status changed to offline.

Action: This is an informational message, and no action is required.

HALRT-01141, Flash Disk was replaced. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8} Flash Cache: {9} Flash Log {10}
Level Warning

Cause: Flash disk was replaced.

Action: This is an informational message, and no action is required.

HALRT-01143, All disk drives have been placed in WriteThrough caching mode. The controller has preserved cache for a disk drive for more than five minutes. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: The preserved cache was not turned off for more than five minutes.

Action: Contact Oracle Support Services.

HALRT-01144, Cell disk was missing and is recreated on the {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Cause: The cell disk was missing and has been recreated on the drive.

Action: See "HALRT-10001".

HALRT-01145, Flash Disk status changed to predictive failure. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8} Flash Cache: {9} Flash Log: {10}
Level Critical

Cause: Flash disk status changed to `predictive failure`.

Action: The flash disk has entered predictive failure status. A white cell locator LED has been lit to help locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait until rebalance and flush have completed before replacing the flash disk.

> **See Also:**
>
> "Maintaining the Flash Disks of Storage Servers".

HALRT-01146, {0} entered poor performance status. The LUN {1} changed status to {2}. Status: {3} Manufacturer: {4} Model Number: {5} Size: {6} Serial Number: {7} Firmware: {8} Slot Number: {9} Cell Disk: {10} Grid Disk: {11} Flash Cache: {12} Flash Log: {13} Reason for poor performance: {14}
Cause: Flash disk status changed to `poor performance`.

Action: See "Maintaining the Flash Disks of Storage Servers".

HALRT-01147, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Warning

Cause: Flash disk status changed to `normal`.

Action: This is an informational message, and no action is required.

HALRT-01148, Flash Disk removed. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8} Flash Cache: {9} Flash Log: {10} Error Count: {11} Last Failure: {12}
Level Critical

Cause: Flash disk was removed.

Action: A flash disk has been removed or is no longer visible to Exadata Storage Server.

- If a flash disk has just been replaced, then this alert is informational only and a flash disk replaced CLEAR alert is sent. No action is needed.
- If flash is being used as a cache, then the only effect is that the flash cache size has been reduced on this cell. The flash disk removal does not cause an outage, data loss, or Oracle ASM rebalance. Each flash card contains four flash disks.

Each flash card contains four flash disks. If all four flash disks on a card alert, then the removal is likely for a card. To see whether the flash disks are on the same card, check for identical PCI Slot numbers in the alert outputs.

HALRT-01149, {0} status changed to offline. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Warning

Cause: Flash disk status changed to `offline`.

Action: This is an informational message, and no action is required.

HALRT-01150, Cell disk was missing and is recreated on the {0}.
Cause: The cell disk was missing, and has been recreated on the disk.

Action: "HALRT-10001".

HALRT-01153, Disk drive {0} is back in WriteBack caching mode.
Cause: The disk drive entered write back caching mode.

Action: See "HALRT-10506".

HALRT-01154, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Cause: The server did not auto-create the cell disk or grid disks.

Action: The storage server software continues to retry, and an alert is sent when the auto-creation succeeds.

HALRT-01155, All disk drives have been temporarily placed in WriteThrough caching mode because the disk controller battery was replaced. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery

Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: Disk controller battery was replaced.

Action: This is an informational message, and no action is required.

HALRT-01156, Previously missing {0} is found. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Clear

Cause: The missing flash disk was found.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then the flash cache are automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then the flash log are automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then the grid disks are automatically recreated and data will be automatically rebalanced by Oracle ASM.

HALRT-01157, Previously missing {0} is found. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Clear

Cause: The missing hard disk was found.

Action: This is an informational message, and no action is required.

HALRT-01159, {0} import failed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Warning

Cause: The import procedure failed.

Action: System failed to import this hard disk because preserved cache exists for other missing hard disks. Reseat or replace the hard disks in {0}. Another alert is sent when all disks are successfully imported.

HALRT-01160, Failed to create LUN on {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Informational

Cause: The disk controller did not create a LUN on the physical disk.

Action: Contact Oracle Support Services.

HALRT-01161, All disk drives are back in WriteBack caching mode. Preserved cache has turned off. Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: Disk drives returned to write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01162, All disk drives are in WriteBack caching mode. Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: Disk drives were in write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01163, All disk drives are back in WriteBack caching mode. Battery has completed charging. Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}.
Cause: Battery completed charging, and all disk drives returned to write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01164, All disk drives are in WriteBack caching mode. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: Disk drives were in write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01165, {0} can be replaced now. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Critical

Cause: Physical disk needed replacement.

Action: Replace the physical disk.

HALRT-01166, {0} can be replaced now. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk {9} Flash Cache {11} Flash Log {12}
Level Critical

Cause: Flash disk needed replacement.

Action: Replace the flash disk.

HALRT-01167, {0} can be replaced now.
Level Critical

Cause: The component needed replacement.

Action: Replace the component.

HALRT-01168, {0} was removed and inserted. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Informational

Cause: The component was inserted.

Action: This is an informational message, and no action is required. CELLSRV may have reported I/O errors when the physical disk was removed. If any grid disk on this physical disk went offline in Oracle ASM as a result, then Oracle ASM automatically brings the disk online.

HALRT-01169, {0} entered write-through caching mode. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Critical

Cause: The component entered write through caching mode.

Action: The flash disk has entered predictive failure status. A white cell locator LED has been lit to help locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait until rebalance and flush have completed before replacing the flash disk.

> **See Also:**
>
> "Maintaining the Flash Disks of Storage Servers".

HALRT-01170, Failed to flush the write back flash cache. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Critical

Cause: Failed to flush the flash cache.

Action: Check the alert log for additional details. After the underlying problem is resolved, flush the write back flash cache on this cell disk is automatically resumed.

HALRT-01171, {0} entered confinement offline status. The LUN {1} changed status to {2}. CellDisk changed status to {3}. All subsequent I/Os on this disk are failed immediately. Confinement tests will be run on the disk to determine if the disk should be dropped. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Reason for confinement: {13}
Cause: The hard disk was performing poorly, and entered confinement.

Action: See "HALRT-10081".

HALRT-01172, {0} entered confinement offline status. The LUN {1} changed status to {15}. CellDisk changed status to {2}. All subsequent I/Os on this disk are failed immediately. Confinement tests will be run on the disk to determine if the disk should be dropped. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Flash Cache: {13} Flash Log: {14} Reason for confinement: {15}
Cause: The flash disk was performing poorly, and entered confinement.

Action: See "HALRT-10082".

HALRT-01173, {0} entered confinement status. The LUN {1} changed status to {2}. CellDisk changed status to {3}. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Reason for confinement: {13}
Cause: The system disk entered confinement.

Action: See "HALRT-10079".

HALRT-01174, {0} entered confinement status. The LUN {1} changed status to {2}. CellDisk changed status to {3}. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Flash Cache: {13} Flash Log: {14} Reason for confinement: {15}
Cause: The disk entered confinement.

Action: See "HALRT-10080".

HALRT-01175, Partition table is missing in system slot {0}.
Level Warning

Cause: The partition table was not on the drive in the system slot.

Action: Contact Oracle Support Services.

HALRT-01176, The data disk from slot {1} was inserted into system slot {2} of the same Oracle Exadata storage server. Status: {3} Manufacturer: {4} Model Number: {5} Size: {6} Serial Number: {7} Firmware: {8} Slot Number: {9}
Level Warning

Cause: The data disk was inserted into the wrong slot.

Action: Insert the disk into the correct slot.

HALRT-01177, The system disk from slot {0} was inserted into data slot {1} of the same Oracle Exadata storage server. Status: {2} Manufacturer: {3} Model Number: {4} Size: {5} Serial Number: {6} Firmware: {7} Slot Number: {8}
Level Warning

Cause: A system disk from the same Exadata Storage Server was inserted in a data disk slot.

Action: Insert the correct disk into the slot.

HALRT-01178, The system disk from a different Oracle Exadata storage server was inserted in slot {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Warning

Cause: A system disk from a different storage server was inserted in the slot.

Action: Insert the correct disk into the slot.

HALRT-01179, A data disk from a different Oracle Exadata storage server was inserted in system slot {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}.
Level Warning

Cause: A data disk from a different storage server was inserted in a system disk slot.

Action: Insert the correct disk into the slot.

HALRT-01181, {0} entered peer failure status. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Critical

Cause: The flash disk entered peer failure status.

Action: The flash disk has entered predictive failure status. A white cell locator LED has been lit to help locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait until rebalance and flush have completed before replacing the flash disk.

HALRT01182, {0} was removed before running the ALTER PHYSICALDISK DROP FOR REPLACEMENT command. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk {9}
Level Critical

Cause: The disk was removed before the `ALTER PHYSICALDISK DROP FOR REPLACEMENT` command was run.

Action: This is an informational message, and no action is required.

HALRT01183, {0} was removed before running the ALTER PHYSICALDISK DROP FOR REPLACEMENT command and inserted. {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Informational

Cause: The disk was removed before the `ALTER PHYSICALDISK DROP FOR REPLACEMENT` command was run.

Action: This is an informational message, and no action is required. If an Oracle ASM grid disk was on the physical disk when the disk went offline, then Oracle ASM automatically brings the disk online.

HALRT-01184, {0} automatically failed over to a spare disk. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Critical

Cause: The cell disk failed over to the spare disk.

Action: This is an informational message, and no action is required.

HALRT-01185, {0} automatically failed over to a spare disk. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Cause: The flash disk failed over to the spare disk.

Action: This is an informational message, and no action is required.

HALRT-01186, {0} was removed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Critical

Cause: A flash disk was removed.

Action: A flash disk was removed or is no longer visible.

- If the flash disk has been replaced, then this alert is informational only, and a flash disk replaced CLEAR alert will be sent. No action is needed.
- If flash is being used as a cache, then the effect is that the flash cache size has reduced on the cell. Flash disk removal does not cause an outage, data loss, or Oracle ASM rebalance.
- Each flash card contains four flash disks. If all four flash disks on a card send an alert, then the removal is likely for a card. To determine if the flash disks are on the same card, check for identical PCI slot numbers in the alert output.

HALRT-01187, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Warning

Cause: The flash disk was replaced.

Action: This is an informational message, and no action is required.

HALRT--01188, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}.Board Tracer Number : {12}
Level Warning

Cause: The flash disk was replaced.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then the flash cache is automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then the flash log is automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then the grid disks are automatically recreated and the data will be automatically rebalanced by Oracle ASM.

HALRT--01189, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Board Tracer Number: {8}
Level Warning

Cause: The component was replaced.

Action: The storage server software continues to retry, and sends an alert when the auto-create operation succeeds.

HALRT--01190, Previously missing {0} is found. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Critical

Cause: The flash disk was located.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then the flash cache is automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then the flash log is automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then the grid disks are automatically recreated, and the data is automatically rebalanced by Oracle ASM.

HALRT--01191, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Warning

Cause: The flash disk status changed to normal.

Action: This is an informational message, and no action is required.

HALRT--01192, {0} was removed and inserted. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Board Tracer Number : {9}
Level Informational

Cause: The physical disk was removed and inserted.

Action: This is an informational message, and no action is required. CELLSRV might have reported I/O errors when the disk was removed. If any grid disk on the physical disk went offline in Oracle ASM, then Oracle ASM automatically brings the disk online.

HALRT--01193, {0} failed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Board Tracer Number : {10}
Level Critical

Cause: A flash disk failed.

Action: Replace the failed flash disk. To identify the affected server, look for a lit white cell locator LED.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.

- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01194, {0} entered predictive failure status. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Critical

Cause: The flash disk entered predictive failure status.

Action: The flash disk has entered predictive failure status. Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01195, {0} entered peer failure status. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number: {12}
Level Critical

Cause: The flash disk entered peer failure status.

Action: The flash disk has entered peer failure status.Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01196, {0} entered write-through caching mode. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Critical

Cause: The flash disk entered write-through caching mode.

Action: The flash disk has entered write-through caching mode. Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01197, {0} sstatus changed to offline. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number: {12}
Level Critical

Cause: The component status changed to offline.

Action: This is an informational message, and no action is required.

HALRT--01198, {0} automatically failed over to a spare disk. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Critical

Cause: The component failed over to a spare disk.

Action: This is an informational message, and no action is required.

Accelerator/Controller Alert Messages: HALRT-01209 to HALRT-01300

HALRT-01209, Alarm has been enabled by user for Adapter {0}.
Cause: An alarm was enabled for the adapter.

Action: See "HALRT-10001".

HALRT-01210, Background initialization rate changed to {0} for Adapter {1}.
Cause: The background initialization rate was changed for the adapter.

Action: See "HALRT-10001".

HALRT-01211, Controller {0} cache discarded due to memory/battery problems.
Cause: The controller cache was discarded due to memory or battery problems.

Action: See "HALRT-10033".

HALRT-01212, Unable to recover Cache Data due to configuration mismatch for Adapter {0}.
Cause: Cache data was recovered due to a configuration mismatch.

Action: See "HALRT-10033".

HALRT-01213, Cache Data Recovered for Adapter {0}.
Cause: Cache data was recovered for the adapter.

Action: See "HALRT-10001".

HALRT-01214, Controller cache discarded due to firmware version incompatibility for Adapter {0}
Cause: Controller cache was discarded because of incompatibility with the firmware version.

Action: See "HALRT-10033".

HALRT-01215, Consistency Check rate changed to {0} for Adapter {1}.
Cause: The consistency check rate was changed for the adapter.

Action: See "HALRT-10001".

HALRT-01216, Flash downloaded image corrupted for Adapter {0}.
Cause: The flash memory card image was corrupted.

Action: See "HALRT-10034".

HALRT-01217, Flash erase error for Adapter {0}.
Cause: There was a flash erase error for the adapter.

Action: See "HALRT-10034".

HALRT-01218, Flash timeout during erase for Adapter {0}.
Cause: The flash timed out during the erase procedure.

Action: See "HALRT-10034".

HALRT-01219, Flash error for Adapter {0}.
Cause: There was a flash error for the adapter.

Action: See "HALRT-10034".

HALRT-01220, Flashing image: {0} for Adapter {1}.
Cause: The controller flash was upgraded with a new image.

Action: See "HALRT-10001".

HALRT-01221, Flash of new firmware image(s) complete for Adapter {0}.
Cause: The new firmware image was copied to flash.

Action: See "HALRT-10001".

HALRT-01222, Flash programming error for Adapter {0}.
Cause: There was a flash programming error for the adapter.

Action: See "HALRT-10034".

HALRT-01223, Flash timeout during programming for Adapter {0}.
Cause: The flash timed out while programming of the adapter.

Action: See "HALRT-10034".

HALRT-01224, Fatal error received for Adapter {0}.
Cause: A fatal error occurred.

Action: See "HALRT-10035".

HALRT-01225, Rebuild rate changed to {0} for Adapter {1}.
Cause: The rebuild rate changed for the adapter.

Action: See "HALRT-10001".

HALRT-01226, Alarm has been disabled by user for Adapter {0}.
Cause: The alarm for the adapter was disabled.

Action: See "HALRT-10001".

HALRT-01227, Configuration cleared for Adapter {0}.
Cause: The configuration for the adapter was cleared.

Action: See "HALRT-10036".

HALRT-01228, Factory defaults restored for Adapter {0}.
Cause: The factory default settings were restored for the adapter.

Action: See "HALRT-10036".

HALRT-01229, HDD disk controller battery not Present for Adapter {0}.
Cause: The battery was not present for the adapter.

Action: See "HALRT-10036".

HALRT-01230, New HDD disk controller battery Detected for Adapter {0}.
Cause: A new battery was detected for the adapter.

Action: See "HALRT-10001".

HALRT-01231, HDD disk controller battery has been replaced for Adapter {0}.
Cause: The battery was replaced for the adapter.

Action: See "HALRT-10001".

HALRT-01232, HDD disk controller battery temperature is high for Adapter {0}.
Cause: The battery temperature was high for the adapter.

Action: See "HALRT-10037".

HALRT-01233, HDD disk controller battery voltage low for Adapter {0}.
Cause: The battery voltage was low for the adapter.

Action: See "HALRT-10038".

HALRT-01234, HDD disk controller battery temperature is normal for Adapter {0}.
Cause: The battery temperature was normal for the adapter.

Action: See "HALRT-10001".

HALRT-01235, HDD disk controller battery needs replacement, SOH Bad for Adapter {0}.
Cause: The battery needed replacement.

Action: See "HALRT-10038".

HALRT-01236, HDD disk controller battery removed for Adapter {0}.
Cause: The battery was removed for the adapter.

Action: See "HALRT-10036".

HALRT-01237, Current capacity of the HDD disk controller battery is below threshold for Adapter {0}.
Cause: The battery capacity was below the threshold for the adapter.

Action: See "HALRT-10038".

HALRT-01238, Current capacity of the HDD disk controller battery is above threshold for Adapter {0}.
Cause: The battery capacity was above the threshold for the adapter.

Action: See "HALRT-10001".

HALRT-01239, Bbu FRU changed for Adapter {0}.
Cause: The battery backup unit (BBU) field replacement unit (FRU) changed for the adapter.

Action: See "HALRT-10001".

HALRT-01240, The disk controller battery is executing a learn cycle and may temporarily enter WriteThrough Caching mode as part of the learn cycle. Disk write throughput might be temporarily lower during this time. The flash drives are not affected. The battery learn cycle is a normal maintenance activity that occurs quarterly and runs for approximately 1 to 12 hours. Note that many learn cycles do not require entering WriteThrough caching mode. When the disk controller cache returns to the normal WriteBack caching mode, an additional informational alert will be sent. Battery Serial Number: {1} Battery Type: {2} Battery Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}
Cause: Disk controller battery started the learn cycle.

Action: This is an informational message, and no action is required. A learn cycle is a normal maintenance activity that occurs quarterly and runs for approximately 1 to 12 hours. The disk controller cache might go into write through caching mode during the learn cycle. Disk write throughput might be temporarily lower during this time.

> **See Also:**
>
> "Monitoring Hard Disk Controller Write-Through Cache Mode"

HALRT-01243, All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. The disk controller battery is executing an unscheduled learn cycle. Battery Serial Number: {1} Battery Type: {2} Battery

Tempearature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperaturer: {6}.
Cause: An unscheduled learn cycle started.

Action: Determine the reason for the unscheduled learn cycle.

HALRT-01244, HDD disk controller SAS lanes have encountered errors.
Cause: A problem was encountered for the hard disk drive disk controller SAS lanes.

Action: See "HALRT-10035".

HALRT-01245, HDD disk controller SAS lanes are back to normal.
Cause: The hard disk drive disk controller SAS lanes were normal.

Action: This is an informational message.

Voltage Alert Messages: HALRT-01301 to HALRT-01400

HALRT-01301, Under-Voltage Problem (lower critical - going low)
Cause: There was an under-voltage problem.

Action: See "HALRT-10001".

HALRT-01302, Under-Voltage Problem Cleared
Cause: The under-voltage problem was cleared.

Action: See "HALRT-10001".

HALRT-01303, Generic Critical Voltage Problem
Cause: There was a generic voltage problem.

Action: See "HALRT-10001".

HALRT-01304, Generic Critical Voltage Problem Cleared
Cause: The generic voltage problem was cleared.

Action: See "HALRT-10001".

Physical Disk Alert Messages: HALRT-01401 to HALRT-01500

HALRT-01401, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Clear

Cause: The component was replaced in the rack.

Action: The grid disks are automatically recreated, and the data is automatically rebalanced by Oracle ASM.

HALRT-01402, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Clear

Cause: The flash disk was replaced in the rack.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then flash cache will be automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then flash log will be automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then grid disks will be automatically recreated and data will be automatically rebalanced by Oracle ASM.

HALRT-01403, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Clear

Cause: The component was replaced in the rack.

Action: This is an informational message, and no action is required.

HALRT-01404, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Clear

Cause: The component was replaced in the rack.

Action: This is an informational message, and no action is required.

HALRT-01405, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Clear

Cause: The component was replaced in the rack.

Action: The storage server software continues to try to create the cell disk and grid disks. An alert is sent when the auto-creation succeeds.

HALRT-01406, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11} Board Tracer Number: {12}
Level Clear

Cause: The component was replaced in the rack.

Action: This is an informational message, and no action is required.

HALRT-01407, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11} Board Tracer Number: {12}
Level Clear

Cause: The flash disk was replaced in the rack.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then flash cache will be automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then flash log will be automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then grid disks will be automatically recreated and data will be automatically rebalanced by Oracle ASM.

HALRT-01408, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Board Tracer Number: {8}
Level Clear

Cause: The component was replaced in the rack.

Action: The storage server software continues to try to create the server disk and grid disks. An alert is sent when the auto-creation succeeds.

HALRT-01409, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Clear

Cause: The component status changed to normal.

Action: This is an informational message, and no action is required.

HALRT-01410, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Clear

Cause: The component status changed to normal.

Action: This is an informational message, and no action is required.

HALRT-01411, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11} Board Tracer Number: {12}
Level Clear

Cause: The component status changed to normal.

Action: This is an informational message, and no action is required.

Auto Service Request Alert Messages for Storage Servers: HALRT-02001 to HALRT-02013

HALRT-02001
Cause: The system disk has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02002
Cause: The system disk had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02003
Cause: A data disk has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02004
Cause: A data disk had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02005
Cause: A flash disk has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02006
Cause: A flash disk had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02007
Cause: A hard disk in a compute server has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02008
Cause: A hard disk in a compute server had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02009
Cause: System hard disk had poor performance.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02010
Cause: Data hard disk had poor performance.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02011
Cause: Flash disk had poor performance.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02012
Cause: The BBU failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02013
Cause: The BBU degraded for unknown reasons.

Action: An alert has been submitted to Oracle Support Services.

HDD Controller Battery Alert Messages: HALRT-02101 to HALRT-02200

HALRT-02101, The HDD disk controller battery was dropped for replacement. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Warning

Cause: The HDD disk controller was dropped before replacement.

Action: Replace the HDD disk controller battery or reenable the HDD disk controller battery using the `ALTER CELL BBU REENABLE` command if the HDD disk controller battery was dropped in error.

HALRT-02102, HDD disk controller battery is {0}. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery failed.

Action: Replace the HDD disk controller battery.

HALRT-02103, All disk drives have been temporarily placed in WriteThrough caching mode because HDD disk controller battery is performing a learn cycle. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Informational

Cause: The HDD disk controller battery was performing a learn cycle.

Action: This is an informational alert. Disk drives will be in `writeback` mode when the HDD disk controller battery completes the learn cycle.

HALRT-02104, All disk drives have been temporarily placed in WriteThrough caching mode because HDD disk controller battery is charging. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Informational

Cause: The HDD disk controller battery was charging.

Action: This is an informational alert. Disk drives will be in `writeback` mode when the HDD disk controller battery completes the charging.

HALRT-02105, The HDD disk controller battery was removed. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected.
Level Informational

Cause: The HDD disk controller battery was removed.

Action: The HDD disk controller battery needs replacement.

HALRT-02106,The HDD disk controller battery capacity was under threshold value of {0} mAh. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery capacity went too low.

Action: Replace the HDD disk controller battery or schedule a learn cycle.

HALRT-02107, The HDD disk controller battery temperature exceeded the threshold value of {6} Celsius. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery temperature was too high.

Action: The HDD disk controller battery was overheated. Check if the ambient temperature is below {0} Celsius and if the fans are working.

HALRT-02108, All disk drives have been temporarily placed in WriteThrough caching mode because HDD disk controller battery was replaced. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Informational

Cause: The HDD disk controller battery was replaced.

Action: This is an informational alert. Disk drives will be in `writeback` mode until the HDD disk controller battery completes charging.

HALRT-02109, All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery entered degraded or unknown status.

Action: Replace the HDD disk controller battery or schedule a learn cycle.

HALRT-02110, The HDD disk controller battery was removed before running the ALTER CELL BBU DROP FOR REPLACEMENT command. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Critical

Cause: The HDD disk controller battery was removed, but the `ALTER CELL BBU DROP FOR REPLACEMENT` command was not run prior to removal of the battery.

Action: Replace the HDD disk controller battery.

HALRT-02111, The HDD disk controller battery was replaced. All disk drives are in WriteBack caching mode. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Clear

Cause: The HDD disk controller battery was replaced.

Action: This is an informational alert and no user action is specified.

HALRT-02112, All disk drives are in WriteBack caching mode. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Clear

Cause: The HDD disk controller battery returned to normal.

Action: This is an informational alert and no user action is specified.

HALRT-02113, All disk drives have been placed in WriteThrough caching mode. The controller has preserved cache for a disk drive for more than five minutes. The flash drives are not affected. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Critical

Cause: All disk drives were placed in `writethrough` caching mode because of preserved cache.

Action: Contact Oracle Support Services.

HALRT-02114, The HDD disk controller battery is performing an unscheduled learn cycle. All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The HDD disk controller battery performed an unscheduled learn cycle.

Action: Determine the cause of the unscheduled learn cycle.

HALRT-02115, Disk drive {0} is in WriteThrough caching mode. Write Performance on this drive will be affected.
Level Critical

Cause: Disk drive {0} entered `writethrough` caching mode.

Action: The HDD disk controller battery is good, but disk drive {0} is in `writethrough` caching mode. Contact Oracle Support Services.

HALRT-02116, Disk drive {0} is back in WriteBack caching mode.
Level Clear

Cause: Disk drive {0} returned to `writeback` caching mode.

Action: This is an informational alert.

HALRT-02117, The HDD disk controller battery is performing a learn cycle. Battery Serial Number: {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity: {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The hard disk controller battery entered a learn cycle.

Action: This is an informational alert, and no user action is needed.
All disk drives may temporarily enter `writethrough` caching mode as part of the learn cycle. Disk write throughput may be lower during this time. The flash drives are not affected.
The battery learn cycle is a normal maintenance activity that occurs quarterly, and runs for approximately 1 to 12 hours. Many learn cycles do not require entering `writethrough` caching mode. When the disk controller cache returns to `writeback` caching mode, an informational alert will be sent.

HALRT-02118, The HDD disk controller battery was replaced. All disk drives are in WriteBack caching mode. Battery Serial Number: {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity: {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The hard disk controller battery was replaced.

Action: This is an informational alert, and no user action is needed.

HALRT-02119, The HDD disk controller battery was replaced. Battery Serial Number: {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity: {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The hard disk controller battery was replaced.

Action: This is an informational alert, and no user action is needed.

Generic Actions for Alert Messages: HALRT-10001 to HALRT-10500

HALRT-10001
Cause: There was a nonerror change for a hardware status, such as a return to normal status.

Action: This is an informational alert and no user action is specified.

HALRT-10003
Cause: An unscheduled learn cycle was started.

Action: Examine how the unscheduled learn cycle was started.

HALRT-10004
Cause: Temperature was not in the acceptable range.

Action: Shut down the system. Determine the cause of thermal failure. If necessary, move the data off of this server.

HALRT-10005
Cause: There was a problem with the fan.

Action: Shut down the system. Determine cause of the fan failure. Replace the fan. If necessary, move the data off of this server.

HALRT-10007
Cause: There was a problem with the power supply.

Action: Examine the power supply. Replace the power supply if necessary.

HALRT-10015
Cause: Temperature was not in the acceptable range.

Action: Examine thermal degradation. Shut down the system and move the data off of this server.

HALRT-10025
Cause: A CPU error was detected.

Action: Shut down the system. Examine CPU error.

HALRT-10029, A generic sensor has detected an error.
Cause: There was a sensor change or error.

Action: Examine the sensor state change or error. If it is unexpected and has not been cleared yet, contact Oracle Support Services.

HALRT-10030
Cause: There was an alert.

Action: Refer to the additional information in the alert message.

HALRT-10031
Cause: The FRU was removed.

Action: If the unit was not intentionally removed, contact Oracle Support Services.

HALRT-10032
Cause: There was an error.

Action: Capture any error message on the console. Restart or power cycle the system and try to bring it back to operation. Contact Oracle Support Services with the error messages.

HALRT-10033
Cause: The cache was lost.

Action: Losing the cache implies that any data on the disks is now stale. Drop force the data on disks attached to this controller from Oracle ASM and then add them back.

HALRT-10034
Cause: The disk controller firmware needed a refresh.

Action: Reflash the disk controller firmware.

HALRT-10035
Cause: The disk controller noticed errors.

Action: Contact Oracle Support Services.

HALRT-10036
Cause: There was a change to the system.

Action: If this change is unintentional, contact Oracle Support Services.

HALRT-10037
Cause: The operating temperature was not within the specified range for the machine.

Action: Ensure the operating temperature is within specified levels for the system.

HALRT-10038
Cause: The battery was not fully charged.

Action: If the battery is not fully charged, then charge it. Otherwise, the battery might be about to fail. Replace disk controller battery.

HALRT-10039
Cause: The LUN on the virtual drive has failed.

Action: Drop the LUN and re-create it.

HALRT-10040
Cause: The physical disk was removed.

Action: If the physical drive was removed unintentionally, replace it immediately.

HALRT-10041
Cause: The disk was not supported or certified for the system.

Action: Replace this disk with a supported and certified disk.

HALRT-10045
Cause: The flash write performance was degraded.

Action: Replace the flash disk, flash PCI card, or energy storage module (ESM). If all disks on the PCI card show degraded performance, then replace the card. Check the ESM LED on the PCI card. The LED is visible from the back of the cell chassis. If this LED is amber, then replace the card.

> **Note:**
>
> This message may include text from message HALRT-10046 if there are grid disks, or from message HALRT-10047 if there are no grid disks but the server disk is used for flash cache.

HALRT-10046
Cause: There was an error with the disk.

Action: If the disks being replaced have persistent data, then rebalance data out of these disks, and replace the flash disk. If there was persistent data on the flash disks, then rebalance data back into the disks.

HALRT-10047
Cause: There was an error with the flash disk.

Action: The flash cache portion provided by the flash disk will be automatically disabled.

HALRT-10054
Cause: The physical disk demonstrated extremely poor performance.

Action: Replace the physical disk at the earliest opportunity.

HALRT-10056
Cause: Auto Service Request was notified with a unique identifier.

Action: This is an informational alert, and no user action is needed.

HALRT-10057
Cause: Detailed information on this problem can be found at My Oracle Support.

Action: Refer to the My Oracle Support note listed in the alert.

HALRT-10066
Cause: The system hard disk demonstrated extremely poor performance.

Action: Replace the system hard disk as soon as possible.

HALRT-10067
Cause: The data hard disk demonstrated extremely poor performance.

Action: Replace the data hard disk as soon as possible.

HALRT-10068
Cause: The flash disk demonstrated extremely poor performance.

Action: Replace the flash disk as soon as possible.

HALRT-10073
Cause: The device monitoring utilities had unexpected output.

Action: A resolved alert will be sent when the problem is resolved. If the problem persists, then Contact Oracle Support Services.

HALRT-10074
Cause: A flash disk was removed or was no longer visible to the storage server.

Action: If the flash disk has just been replaced, then this alert is informational only. If the flash disk is being used as a cache, then the flash cache size has been reduced on this server. Flash disk removal will not cause an outage, data loss, or Oracle ASM rebalance.

HALRT-10075
Cause: The temperature of Recovery Appliance is outside the operating range.

Action: Recovery Appliance has an operating range of 5 to 32 degrees Celsius (41 to 89.6 degrees Fahrenheit). Adjust the data center conditions so that Recovery Appliance is operating within the specified temperature range.

HALRT-10076
Cause: The Oracle ASM disks associated with the grid disks on the physical disk were successfully dropped, and a blue "OK to Remove LED" was lit on the disk to indicate it is safe to remove the disk.

Action: Replace the disk drive at the earliest opportunity.

HALRT-10078
Cause: CELLSRV reported I/O errors when the physical disk was removed.

Action: No action is needed. If a grid disk on this physical disk went offline in Oracle ASM as a result, then Oracle ASM will automatically bring the disk online.

HALRT-10079
Cause: The single system disk in the storage server did not have valid partitions.

Action: Contact Oracle Support Services.

HALRT-10080
Cause: A disk was inserted in the wrong slot.

Action: Insert the disk into the correct slot.

HALRT-10081
Cause: The hard disk demonstrated temporary poor performance. The disk has been confined and is in an offline state.

Action: If the hard disk was used for grid disks, then the Oracle ASM disks associated with the grid disks on this hard disk are offline. Confinement tests run on this disk, and either move the disk permanently to `poor performance` or return the disk back to `normal` state.

HALRT-10082
Cause: The flash disk demonstrated temporary poor performance. The disk has been confined and is in an offline state.

Action: One of the following actions is performed:

- If the flash disk was used for Write-Through flash cache, then flash cache on this disk will be disabled causing a reduction in the effective flash cache size.
- If the flash disk was used for Write-Back flash cache, then flash cache on this disk will be disabled, and read I/O to dirty data in the flash cache will go to the mirror disks on different cells. Write I/Os will go directly to the grid disks on this cell.
- If the flash disk was used for flash log, then flash log is disabled on this disk causing a reduction of the effective flash log size.
- If the flash disk was used for `writethrough` flash cache, then flash cache on this disk will be disabled causing a reduction in the effective flash cache size.

- If the flash disk was used for `writeback` flash cache, then flash cache on this disk will be disabled, and read I/O to dirty data in the flash cache will go to the mirror disks on different cells. Write I/Os will go directly to the grid disks on this cell.
- If the flash disk was used for flash log, then flash log is disabled on this disk causing a reduction of the effective flash log size.
- If the flash disk was used for grid disks, then the Oracle ASM disks associated with the grid disks on this flash disk are offline.

Confinement tests run on this disk, and either move the disk permanently to `poor performance` or return the disk back to `normal` state.

HALRT-10086
Cause: The flash disk has entered `peer failure` status.

Action: Another alert will be sent when all data has been rebalanced and flushed from the flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected server.

HALRT-10087
Cause: The data on this disk has been successfully rebalanced and flushed.

Action: Replace the drive.

Generic Actions for Alert Messages: HALRT-10503 to HALRT-11000

These messages are user actions for alert messages. The actions are in the range HALRT-10503 through HALRT-11000.

HALRT-10503
Cause: The HDD disk controller had a problem.

Action: HDD disk controller battery needs replacement. Contact Oracle Support Services.

HALRT-10504
Cause: The battery was in write through caching mode.

Action: HDD disk controller battery is good, but the logical drive is in write through caching mode. Contact Oracle Support

HALRT-10505
Cause: The disk controller battery completed the learn cycle.

Action: Disk controller battery is back to a good state, but the disk drive is in `writethrough` caching mode. Contact Oracle Support Services.

HALRT-10506
Cause: Disk drive entered `writeback` caching mode.

Action: Disk drive is in `writeback` caching mode. No action is needed.

HALRT-10507
Cause: All logical drives were in write through caching mode.

Action: The disk controller battery is good, but all logical drives are in write-through caching mode. Contact Oracle Support Services.

HALRT-10510
Cause: The contents of the disk controller cache have been lost.

Action: Contact Oracle Support Services.
The loss of the disk controller cache potentially leads to some loss of recent write operations. Access to the grid disks on this server has been disabled. The storage server software disk auto management has been disabled for the storage server, to avoid potential conflicts with user-initiated disk offline or drop force during the recovery period. Auto management is not affected on the servers that do not have this problem.

HALRT-10511
Cause: New battery was being charged.

Action: This is an informational message, and no action is required. The new battery is being charged automatically. When the battery completes charging, the disk drives will be in writeback caching mode.

HALRT-10512
Cause: Disk controller battery overheated.

Action: Disk controller battery is over-heated. Check if the ambient temperature is below 32 degrees Celsius and if the fans are working.

HALRT-10513
Cause: Disk controller battery was not functioning.

Action: Replace the disk controller battery.

Generic Actions for Alert Messages: HALRT-11001 to HALRT-11100

HALRT-11001
Cause: There was a problem with the InfiniBand configuration.

Action: Use `ibstatus` to check the InfiniBand configuration, and verify that both ports show state ACTIVE and the expected Gb per second. The alert is cleared when the port status is ACTIVE.
A common cause of the InfiniBand link negotiating to a lower link rate is that the InfiniBand cable connectors on the HCA or the switch are loose. Tightening the InfiniBand connector in the HCA and the switch usually solves this problem.
A port may show status DOWN after a switch restart. The port is automatically made ACTIVE after a few minutes.

> **See Also:**
>
> My Oracle Support note 1612884.1 for additional information

HALRT-11002
Cause: There was a problem with the InfiniBand port.

Action: Use the CellCLI `LIST IBPORT DETAIL` command to check the port. Verify both ports show zero error counts for `rcvErrs` and `symbolErrs`. A non-zero error count can be caused by a bad cable or a cable that is not properly seated. It can also be caused by a bad port on the HCA or switch. Manually disabling and enabling the InfiniBand ports or resetting the InfiniBand switches can also lead to these errors. Correct the problem, and then reset the InfiniBand error counts on both the HCA port and the peer switch port by using the CellCLI `ALTER IBPORT...RESET` command.

HALRT-11003
Cause: There was a problem with the InfiniBand port.

Action: Use the CellCLI `LIST IBPORT DETAIL` command to check the port. Verify that both ports show the expected Gb per second. A common cause for the InfiniBand link negotiating to a lower link rate is that the InfiniBand cable connectors on the HCA or the switch are loose. Tightening the InfiniBand connector in the HCA and in the switch usually solves this problem.

C

Replacement Units

This appendix lists the replacement units for Recovery Appliance. It contains the following topics:

About Warranty Coverage

All replaceable parts are field replaceable units (FRUs) under the Oracle Premier Support for Systems warranty. Oracle field engineers install the replacement parts.

Understanding Repair Categories

Repair categories describe the potential impact of a repair action. Table C-1 describes the categories and actions.

Table C-1 Recovery Appliance Repair Categories

Category	Description	Examples
Hot Swap (HS)	Repair part is hot swappable, and can be replaced without shutting down the host system. Commands might be required before and after replacement to protect data.	• Disks • Fans • Power supplies
Infrastructure Repair (IR)	Repairing a connectivity component within the Recovery Appliance rack. No downtime of the rack is required; however, individual components may require downtime.	• Ethernet switch • External cables • InfiniBand switch

Table C-1 (Cont.) Recovery Appliance Repair Categories

Category	Description	Examples
Compute Server Offline (DBO)	Repairing the part requires shutting down a compute server. No downtime of the rack is required; however, individual servers might require downtime, and they might be taken outside of the cluster temporarily. If the server is currently running, then you must shut it down. See "Powering Down the Servers". If you monitor the system using Oracle Enterprise Manager Grid Control, then avoid unnecessary alerts by putting the target in a blackout state. Refer to *Oracle Enterprise Manager Administration* for additional information.	• Memory • PCIe cards • PDUs • Processors • System boards
Storage server Offline (EBO)	Repairing the part requires shutting down a storage server within the rack. No downtime of the rack is required; however, individual servers might require downtime, and they might be taken outside of the cluster temporarily. There also might be a temporary performance impact, while the system rebalances to maintain data redundancy. If the server is currently running, then you must shut it down. See "Powering Down the Servers". If you monitor the system using Oracle Enterprise Manager Grid Control, then avoid unnecessary alerts by putting the target in a blackout state. Refer to *Oracle Enterprise Manager Administration* for additional information.	• Memory • PCIe cards • Processors • System boards

LED Status Descriptions

This section describes the LEDs.

Oracle Server LEDs

Table C-2 identifies the LED status indicators that can indicate faults on the compute servers and the storage servers.

Table C-2 Oracle Server LED Status Descriptions

Component	LED Status
General Status	The following LEDs are located on the front and back panels of each server: **Service Required (amber)** Off indicates normal operation. Steady on indicates a fault in the server. **Power/OK (green)** Steady on indicates that the server is powered on. A slow blink indicates that the server is starting up. A steady blink indicates that standby power is on, but chassis power is off. Off indicates no AC power or an incomplete startup of Oracle ILOM.
Fans	Each fan module has one bicolored LED, which you can see when the top cover fan door is open. **Fan Status (amber/green)** Amber indicates that the fan module is faulty. The front and rear panel Service Required LEDs are also lit when the system detects a fan module fault. Green indicates that the system is powered on, and the fan module is functioning correctly.
Drives	The following LEDs are located on the storage and boot drives: **Service Required (amber)** Off indicates normal operation. Steady on indicates a fault in the drive. **OK/Activity (green)** Steady on indicates normal inactivity. A steady blink indicates normal activity. Off indicates that the power is off, or the system does not recognize the drive.

Table C-2 (Cont.) Oracle Server LED Status Descriptions

Component	LED Status
Power supply	The following LEDs are on each power supply. You can see them from the back of the server. **Service Required (amber)** Off indicates normal operation. Steady on indicates that the system has detected a problem in the power supply regarding the voltage, current, temperature, or fans. **AC OK/DC OK (green)** Steady on indicates normal operation for both AC and DC power. A slow blink indicates that input AC power is normal and output DC power is disabled. Off indicates no input AC power.

See Also:

- "About Server and Component Status Indicators" in the *Sun Server X4-2 Service Manual* at

  `http://docs.oracle.com/cd/E36975_01/html/E38045/z4000be71405942.html#scrolltoc`

- "About Server and Component Status Indicators" in the *Sun Server X4-2L Service Manual* at

  `http://docs.oracle.com/cd/E36974_01/html/E38145/z40003ec1405942.html#scrolltoc`

Sun Datacenter InfiniBand Switch 36 LEDs

Table C-3 identifies the LED status indicators that can indicate faults on Sun Datacenter InfiniBand Switch 36.

Table C-3 Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Chassis	The following LEDs are located on the left side of the rear panel. **Attention (amber)** Steady on indicates a fault. Off indicates normal operation. Flashing indicates no function. **OK (green)** Steady on indicates normal operation. Off indicates that the switch is off or initializing. Flashing indicates no function.
Link status	The link status LEDs are located by the data cable connectors on the rear panel. **Link (green)** Steady on indicates that a link is established. Off indicates that the link is down. Flashing indicates that there are symbol errors.
Network management ports	The following LEDs are located on the network management connector on the rear panel. **Link Speed (amber/green):** Green steady on indicates a 1000BASE-T link. Amber steady on indicates a 100BASE-T link. Off indicates no link. Flashing indicates no function. **Activity (green):** Flashing indicates packet activity. On indicates no function. Off indicates no activity.

Table C-3 (Cont.) Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Power supply	The following LEDs are located on the power supply at the front of the chassis. **Attention (amber)** Steady on indicates a fault, and 12 VDC is shut down. Off indicates normal operation. Flashing indicates no function. **OK (green)** Steady on indicates 12 VDC. Off indicates no DC voltage. Flashing indicates no function. **AC LED is green** Steady on indicates AC power is supplied. Off indicates that AC power is not present. Flashing indicates no function.
Fan	The following LED is located on the air inlet on the front of the chassis. **FT** On indicates a problem with the fan. Off indicates normal operation

See Also:

"Identifying LEDs" in the *Sun Datacenter InfiniBand Switch 36 Service Manual* at

`http://docs.oracle.com/cd/E26698_01/html/E26434/z40000fe1394411.html`

Parts for Compute Servers

The following topics identify the replacement parts for the Recovery Appliance compute servers:

- Database Server Components of Recovery Appliance X9M-2
- Database Server Components of Recovery Appliance X8M-2
- Oracle Server X8-2 and X8M-2 Database Server Parts

Database Server Components of Recovery Appliance X9M-2

Each Recovery Appliance X9M-2 database server includes the following components:

- CPU - 32-core Intel Xeon 8358 processors (2.6 GHz):
  - Flexible configuration: 2 CPUs
  - Eighth Rack: 1 CPU
- RAM:
  - Flexible configuration: 512 GB RAM (16 x 32 GB DIMMs), expandable to 1 TB (16 x 64 GB DIMMs), 1.5 TB (16 x 32 GB and 16 x 64 GB DIMMs), or 2 TB (32 x 64 GB DIMMs) with memory expansion kit
  - Eighth Rack: 384 GB RAM (12 x 32 GB DIMMs), expandable to 1 TB (16 x 64 GB DIMMs) with memory expansion kit
- System storage: 2 x 3.84 TB PCIe 4.0 NVMe storage drives, expandable to 4 drives
- RDMA Network Fabric: 1 x dual-port CX5 RDMA Network Fabric card; PCIe 4.0, 200 Gb/s aggregate throughput (active-active)
- Administration network: 1 x 1 GbE Base-T Ethernet port
- Integrated Lights Out Manager (ILOM) port for remote management: 1 x 1 GbE Base-T Ethernet port
- Client and additional network connectivity options:
  - Flexible configuration options (one of following):
    - 4 x SFP+/SFP28 ports (10/25 GbE)
    - 8 x 10GBASE-T ports (10 GbE)
    - 2 x SFP+/SFP28 ports (10/25 GbE) and 4 x 10GBASE-T ports (10 GbE)
    - 6 x SFP+/SFP28 ports (10/25 GbE)
    - 4 x SFP+/SFP28 ports (10/25 GbE) and 4 x 10GBASE-T ports (10 GbE)
    - 2 x SFP+/SFP28 ports (10/25 GbE) and 8 x 10GBASE-T ports (10 GbE)
    - 12 x 10GBASE-T ports (10 GbE)
  - Eighth Rack options (one of following):
    - 4 x SFP+/SFP28 ports (10/25 GbE)
    - 8 x 10GBASE-T ports (10 GbE)
    - 2 x SFP+/SFP28 ports (10/25 GbE) and 4 x 10GBASE-T ports (10 GbE)

Database Server Components of Recovery Appliance X8M-2

Recovery Appliance X8M-2 database servers include the following components:

- 2 x 24-core Intel Xeon Platinum 8260 Processors (2.4 GHz)
- 384 GB RAM (12 x 32 GB DIMMs), expandable to 1.5 TB (24 x 64 GB DIMMs) with memory expansion kit
- 4 x 1.2 TB boot drive/hard disks, hot swappable, expandable to 8 drives

- Disk controller HBA with 2 GB cache (no batteries)
- RDMA Network Fabric: 1 x dual-port QSFP28 100 Gb/s PCIe 3.0 RDMA Network Fabric card - both ports active
- Client network: 2 x SFP+/SFP28 ports (10/25 GbE) or 2 x 10GBASE-T ports (10 GbE)
- Backup/optional networks: 2 x SFP+/SFP28 ports (10/25 GbE) or 4 x 10GBASE-T ports (10 GbE)
- Administration network: 1 x 1 GbE Base-T Ethernet port
- Integrated Lights Out Manager (ILOM) port for remote management: 1 x 1 GbE Base-T Ethernet port

Oracle Server X8-2 and X8M-2 Database Server Parts

The following image shows the front view of Oracle Server X8-2 and X8M-2 Database Servers.

Figure C-1 Front View of Oracle Server X8-2 and X8M-2 Database Servers

The following image shows the rear view of the Oracle Server.

Figure C-2 Rear View of Oracle Server X8-2 and X8M-2 Database Servers

The following table lists the replaceable parts for Oracle Server X8-2 and X8M-2 database servers. Refer to this table only if you have a basic warranty. All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table C-4 Replacement Parts for Oracle Server X8-2 and X8M-2 Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
370-5621-01	3 volt lithium battery grover	DBO	FRU	See "Servicing the Battery (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqviw.html#scrolltoc

Table C-4 (Cont.) Replacement Parts for Oracle Server X8-2 and X8M-2 Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7363537	ASSY DRV 1.2TB 10KRPM SAS3 512N SFF MARLIN	DBO	FRU	See "Servicing Storage Drives (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gquak.html#scrolltoc
7322171	ASM LFIM	DBO	FRU	See "Servicing Other Cables (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvev.html#scrolltoc
7339763	BD LP DUAL SFP28 ADAPTER WHITNEY+ PCIe3.0x8	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc
8200663	HEATSINK CPU HEIGHT:23.5 MM	DBO	FRU	See "Servicing Processors (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqtmk.html#scrolltoc
7343599	Assembly, riser, 1-slot, 1U	DBO	FRU	See "Servicing PCIe Risers (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvft.html#scrolltoc
7330698	DIMM, 32 GB, RDIMM, DDR4, 2666 2Rx4 1.2V	DBO	CRU	See "Servicing the DIMMs (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvkr.html#scrolltoc
8200116	Power supply, AC, A266, F, 12V, 1200W	HS	CRU	1. Remove rear left side of the CMA, and then swing it back and out to the right. The telescopic cross member supporting the cables directly behind the server can be retracted to access both power supplies. 2. See "Servicing Power Supplies (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqunc.html#scrolltoc 3. Replace the CMA.

Table C-4 (Cont.) Replacement Parts for Oracle Server X8-2 and X8M-2 Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7332895	HBA, 12 Gb SAS PCIe 16-port RAID internal with flash module	DBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvmq.html#scrolltoc
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	DBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvmq.html#scrolltoc
7340897	Cable, SAS3 to HBA	DBO	CRU	See "Servicing the Internal HBA Cables" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvav.html#scrolltoc
7331123	Assembly, fan module	HS	CRU	See "Servicing Fan Modules (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gquhg.html#scrolltoc
8200428	CPU INTEL CLX P8260 (B1) 2.4 GHz 165W 24c 8S 1TB (S-RF9H)	DBO	FRU	See "Servicing Processors (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqtmk.html#scrolltoc
7104073	Assembly, Dual 40Gb/Sec (4x) QDR RDMA Fabric Host Channel Adapter Module M3	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc 3. See "Configuring the InfiniBand Switches"
7046442 (X8-2)	Assembly, Dual port, QDR CX3, HCA	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc 3. See "Configuring the InfiniBand Switches"

Table C-4 (Cont.) Replacement Parts for Oracle Server X8-2 and X8M-2 Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7119792 (X8M-2)	CX5 100GbE Adapter Dual Port QSFP28	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc 3. See Verifying the RoCE Network Fabric Configuration.
7360829	Assembly, motherboard, 1U	DBO	FRU	1. See "Servicing the Motherboard Assembly (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvfa.html#scrolltoc 2. Press F2 after restarting the server. 3. Set the boot order.
7331100	Assembly, riser, 2-slot, 1U	DBO	FRU	See "Servicing PCIe Risers (CRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvft.html#scrolltoc
7331142	Assembly, Front LED indicator module	DBO	FRU	See "Servicing the Front LED Indicator Module (FRU)" in *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvok.html#scrolltoc

Parts for Storage Servers

The following topics identify the replacement parts for the Recovery Appliance storage servers:

- Oracle Exadata Storage Server X9M-2, X8M-2, and X8-2 High Capacity and Extended (XT) Server Parts
- #unique_455
- #unique_456
- #unique_457
- #unique_458

Oracle Exadata Storage Server X9M-2, X8M-2, and X8-2 High Capacity and Extended (XT) Server Parts

The following figure shows the front view of the Oracle Exadata Storage Server X9M-2, X8M-2 and X8-2 High Capacity and XT servers.

> **Note:**
>
> The Recovery Appliance supports Oracle Exadata Storage Server high capacity servers and does not use Oracle Exadata Storage Server XT servers. Both are discussed for completeness and compatibility with Oracle Exadata documentation.

Figure C-3 Front View of Oracle Exadata Storage Server X9M-2, X8M-2 and X8-2 High Capacity and XT servers

The following figure shows the rear view of the Oracle Exadata Storage Server X8M-2 and X8-2 High Capacity and XT servers.

> **Note:**
>
> The Recovery Appliance does not support XT servers.

Figure C-4 Rear View of Oracle Exadata Storage Server X9M-2, X8M-2 and X8-2 High Capacity and XT servers

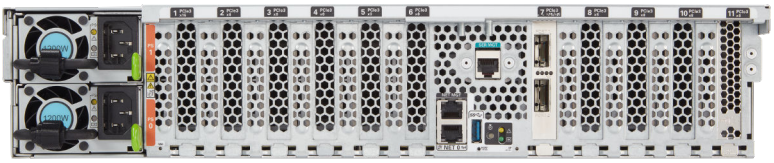

The following table lists the replaceable parts for Oracle Exadata Storage Server X9M-2, X8M-2 and X8-2 High Capacity and XT servers. Refer to this table only if you have a basic warranty. All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

> **Note:**
>
> The Recovery Appliance supports Oracle Exadata Storage Server high capacity servers and does not use Oracle Exadata Storage Server XT servers. Both are discussed for completeness and compatibility with Oracle Exadata.

Table C-5 Replacement Parts for Oracle Exadata Storage Server X8M-2 and X8-2 High Capacity and XT Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
370-5621-01	3 volt lithium battery grover	EBO	FRU	See "Servicing the Battery" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtff.html#scrolltoc
7322171	Assembly, Front LED Indicator Module (LFIM)	EBO	FRU	See "Servicing the Front LED Indicator Module" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtpg.html#scrolltoc
7341141	Assembly, 12-slot disk backplane	EBO	FRU	See "Servicing the Disk Backplane" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtsb.html#scrolltoc
8200116	Power Supply, A266 F, 1200W, AC Input	HS	FRU	1. Remove rear left side of the CMA, and then swing it back and out to the right. The telescopic cross member supporting the cables directly behind the server can be retracted to access both power supplies. 2. See "Servicing Power Supplies" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtag.html#scrolltoc 3. Replace the CMA.
7322166	Fan assembly, 2U	HS	FRU	See "Servicing Fan Modules" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtak.html#scrolltoc

Table C-5 (Cont.) Replacement Parts for Oracle Exadata Storage Server X8M-2 and X8-2 High Capacity and XT Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7046442 (X8-2)	Assembly, dual port, QDR CX3, Host Channel Adapter	EBO	FRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcq.html#scrolltoc 3. See "Configuring the InfiniBand Switches" 4. Verify the new hardware meets the hardware and firmware specifications using the following command: `/opt/oracle.SupportTools/CheckHWnFWprofile`
7119792 (X8M-8)	CX5 100GbE Adapter Dual Port QSFP28	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcq.html#scrolltoc 3. See "Configuring the InfiniBand Switches" 4. Verify the new hardware meets the hardware and firmware specifications using the following command: `/opt/oracle.SupportTools/CheckHWnFWprofile`
7322160	Heat sink, CPU	EBO	FRU	See "Servicing Processors" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtmk.html#scrolltoc
7357761	Assembly, Drive, 14TB 3.5" SAS3/7200 rpm Coral with encryption	HS	FRU	See "Servicing Storage Drives" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcs.html#scrolltoc

Table C-5 (Cont.) Replacement Parts for Oracle Exadata Storage Server X8M-2 and X8-2 High Capacity and XT Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7353916	RDIMM 16GB, DDR4 2666, 1Rx4 1.2V	EBO	FRU	See "Servicing the DIMMs" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcm.html#scrolltoc
7332895	HBA, 12 Gbps SAS PCIe 16-port internal RAID with flash module	EBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtjv.html#scrolltoc
7086345	Assembly, super cap, 13.5V, 6.4F, SAS3 HBA	EBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtjv.html#scrolltoc
7315464	SAS-3 Data Cable Assembly	EBO	FRU	See "Servicing SAS Cables" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqthb.html#scrolltoc
7361253	M.2,240GB,SATA,22x 80mm,SSD,INTEL, DC S4510	EBO	FRU	See "Servicing the Internal M.2 Flash SSDs" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/grlly.html#scrolltoc
7335943	6.4TB FLASH F640 NVMe AIC AURA 7	EBO	FRU	See "Servicing PCIe Cards" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcq.html#scrolltoc
7361454	6.4TB FLASH F640 V2 NVMe AIC AURA 8	EBO	FRU	See "Servicing PCIe Cards" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcq.html#scrolltoc
8200426	CPU INTEL CLX G5218 (B1) 2.3 GHz 125W 16c 4S 1TB (S-RF8T)	EBO	FRU	See "Servicing Processors" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtmk.html#scrolltoc
7360978	Assembly, System board	EBO	FRU	See "Servicing the Motherboard Assembly" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtwc.html#scrolltoc
7341466	Assembly, M.2 form factor riser	EBO	FRU	See "Servicing the Internal M.2 Flash SSDs" in *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at http://docs.oracle.com/cd/E93361_01/html/E93395/grlly.html#scrolltoc

Parts for Cisco Switch

Table C-6 lists replaceable parts for the Cisco switch used with Recovery Appliance.

> **Note:**
>
> All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table C-6 Replacement Parts for Cisco Switch

Part Number	Description	Repair Category	Repair
7023685	Cisco Catalyst 4948 switch, RoHS:Y	IR	See chapter 3 in *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
7024423	Power supply for Cisco Catalyst 4948 switch, RoHS:Y	HS	See chapter 4 in *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
7024424	Cooling fan for Cisco Catalyst 4948 switch, RoHS:Y	HS	See chapter 4 in *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Parts for InfiniBand Switches

Table C-7 lists the replaceable parts for the InfiniBand switches used in Recovery Appliance. See "Cables and Cable Accessories for Recovery Appliance" for information about InfiniBand cables.

> **Note:**
>
> All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table C-7 Replacement Parts for InfiniBand Switches

Part Number	Description	Repair Category	Repair
300-2143	760 watt power supply, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
300-2299	760 watt AC input power supply	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
350-1312	Fan module for Sun Datacenter InfiniBand Switch 36, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
350-1566	Fan module, 1 RU, Sanace, reverse	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
371-2210	CR2032 3 v battery, RoHS:Y	IR	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
541-3495	Sun Datacenter InfiniBand Switch 36, RoHS:YL	IR	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
7052970	Sun Datacenter InfiniBand Switch 36, RoHS:YL	IR	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
7061031	Fan module, 1 RU, Sanace, reverse	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html

Parts for Power Distribution Units

Table C-8 lists the replacement parts for the power distribution units (PDUs). Given the number of cables and connections, Oracle recommends that you shut down the Recovery Appliance rack when replacing power distribution unit parts.

> **Note:**
>
> All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table C-8 Replacement Parts for Power Distribution Units

Part Number	Description	Repair Category	Repair
180-2379	13A/250V 90 degree IEC 320 Sheet E to IEC 320 C13 cord, 1 meter, RoHS:Y	HS	Not applicable
180-2380	13A/250V 90 degree IEC 320 Sheet E to IEC 320 C13 cord, 2 meter, RoHS:Y	HS	Not applicable
371-4000	Three-phase PDU: 15kV,5-Pin, IEC309 32A, 5 PIN 230/400V, 3ph IP44, International, RoHS:Y	IR	See "Replacing a Power Distribution Unit" of this guide.
371-4199	Single-phase PDU: 15kVA, with three 30A/250V 2-Pole/3-Wire NEMA L6-30P Plugs, North America, RoHS:Y	IR	See "Replacing a Power Distribution Unit" of this guide.
371-4201	Single-phase PDU: 15kVA, with three Blue 32A/240V Splash-proof 2-Pole/3-Wire IEC 60309 Plugs, International, RoHS:Y	IR	See "Replacing a Power Distribution Unit" of this guide.
371-5039	Single-phase PDU: 22k VA Hubbell CS8265C PlugNo X-Option	IR	See "Replacing a Power Distribution Unit" of this guide.
371-5040	Single-phase PDU: 22k VA Blue 32A IEC 60309 Plug No X-Option	IR	See "Replacing a Power Distribution Unit" of this guide.
7051908	Three-phase PDU: 15kVA, IEC309 60A, 4 PIN, 250VAC, 3ph IP67, North America, RoHS:Y	IR	See "Replacing a Power Distribution Unit" of this guide.

Cables and Cable Accessories for Recovery Appliance

Table C-9 lists the replacement cables and cable accessories for Recovery Appliance.

> **Note:**
>
> All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table C-9 Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
350-1287	Slim rail, CMA, 1U-2U for a storage server	EBO	1. If the server is currently running, then it should be shut down gracefully. See "Powering On and Off a Recovery Appliance Rack" for additional information. 2. See "Installing the Cable Management Arm" in *Sun Fire X4170, X4270, and X4275 Servers Installation Manual* at http://docs.oracle.com/cd/E19477-01/820-5827-13/rack_install.html#50614309_99582
350-1546	Cable management arm (CMA) 1U, RoHS:Y for compute server	DBO	1. If the server is currently running, then it should be shut down gracefully. See "Powering On and Off a Recovery Appliance Rack" for additional information. 2. See "Installing the Cable Management Arm" in *Sun Fire X4170, X4270, and X4275 Servers Installation Manual* at http://docs.oracle.com/cd/E19477-01/820-5827-13/rack_install.html#50614309_99582
350-1519	Serial cable kit, RoHS:Y with the following: • USB to DB 9-pin M serial cable • DB 9-pin F to DB 9-pin F null modem cable	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
350-1637	Label assembly, EV2, 2-way, upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1638	Label assembly EV2, 3-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1639	Label assembly EV2, 4-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1640	Label assembly EV2, 5-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1641	Label assembly EV2, 6-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1642	Label assembly EV2, 7-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1643	Label assembly EV2, 8-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4

Table C-9 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
530-4445	QSFP InfiniBand passive copper cable, 10 GbE, 3 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: There is a loss of redundancy when replacing cables.
530-4446	QSFP InfiniBand passive copper cable, 10 GbE, 5 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: There is a loss of redundancy when replacing cables.
530-4403	4x QSFP copper QDR InfiniBand cable, 2 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: There is a loss of redundancy when replacing cables.
530-4404	4x QSFP copper QDR InfiniBand cable, 3 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4415	4x QSFP copper QDR InfiniBand cable, 5 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4432	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, blue, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4433	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, red, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4434	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, red, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4435	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, black, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Table C-9 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
530-4436	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, black, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4437	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, orange, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4438	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, black, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4444	QSFP to QSFP Passive Copper Cable, 1 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4445	QSFP passive copper cable, 10 GE, 3 m	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4502-01	Ethernet cable, Cat 5/5E, 10 foot, blue, RoHS:Y	HS	See *Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
7015400	Power cable, jumper, SR2, 2m C14RA, 10A, C13	HS	See "Replacing a Power Distribution Unit" of this guide.
7042273	Slim rail, cable management arm, 1U-2U for a storage server	EBO	1. If the server is currently running, then shut it down. See "Shutting Down Recovery Appliance". 2. See "Installing the Second-Generation Cable Management Arm" in *Sun Server X4-2L Service Manual* at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000151427029.html#scrolltoc
7102869	Passive fiber InfiniBand cable, 10 m, RoHS:Y	HS	1. See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html **Note:** Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.

Table C-9 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
7102870	Passive fiber InfiniBand cable, 20 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html **Note:** Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.
7102871	Passive fiber InfiniBand cable, 50 m, RoHS:Y	HS	See *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html **Note:** Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.

D

In-Rack Cabling Tables for Recovery Appliance X9M

This appendix identifies the location of the various components in a Recovery Appliance X9M full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5

This appendix identifies the location of the various components in a Recovery Appliance X8, X7, X6, and X5 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

This appendix contains the following sections:

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black**: RDMA Network Fabric cables, AC power jumper cables, and Ethernet cables for storage servers
- **Blue**: Gigabit Ethernet cables for compute servers
- **Red**: Integrated Lights Out Manager (ILOM) cables

Recovery Appliance X8, X7, X6, and X5 Rack Layout

Figure D-1 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 full rack. The front is shown on the left, and the rear is shown on the right.

Note:

The 1U empty space at the bottom of the rack is reserved for an optional InfiniBand spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure D-1 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Full Rack

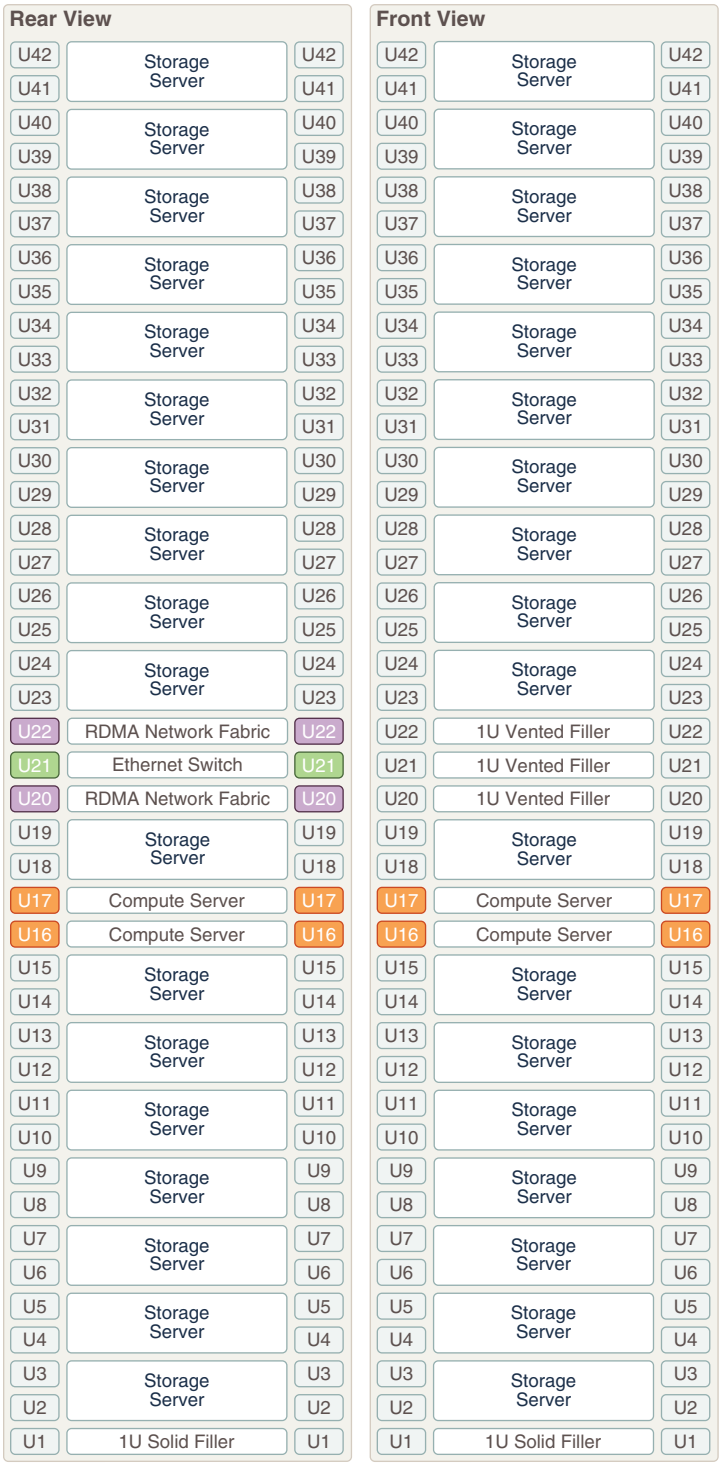

Figure D-2 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 minimum configuration.

Figure D-2 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Minimum Configuration

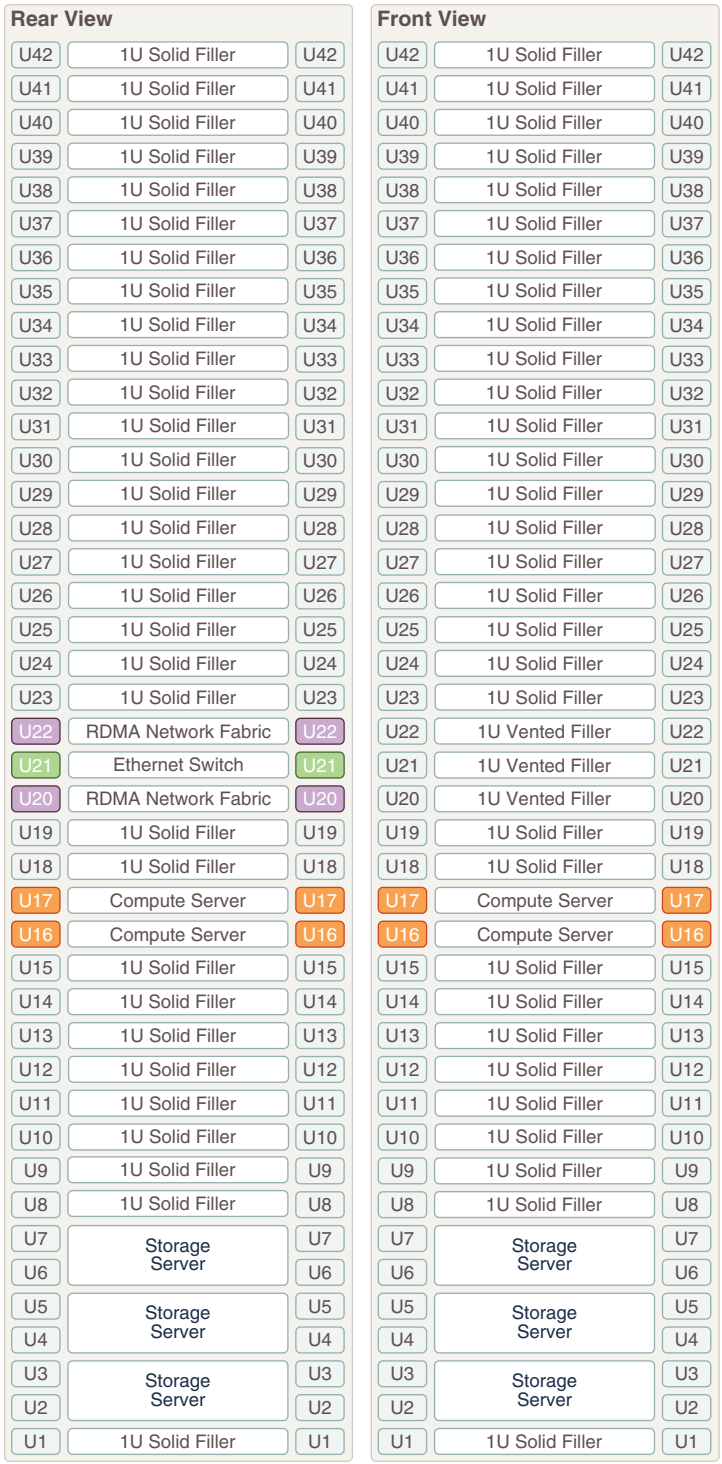

Oracle Integrated Lights Out Manager Cabling Connections

Table D-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in

racks with fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table D-1 ILOM Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	16
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U25	Storage server	18
U23	Storage server	22
U18	Storage server	26
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

1 U*n* is the unit location in the rack, where *n* is the number.

Administrative Gigabit Ethernet Cabling Connections

Table D-2 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in racks with fewer storage servers. The Ethernet switch is located in rack unit 21. Ethernet cables are blue.

Table D-2 Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	15
U39	Storage server	1
U37	Storage server	3

Table D-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U35	Storage server	5
U33	Storage server	7
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U25	Storage server	17
U23	Storage server	21
U22	InfiniBand leaf switch	45
U20	InfiniBand leaf switch	46
U18	Storage server	25
U17	Compute server	27
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X8

Table D-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-3 Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G3-6	G2-0	2 meters
U39	G5-4	G0-2	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-2	G0-4	2 meters
U33	G5-0	G0-6	2 meters
U31	G4-2	G1-4	2 meters

Table D-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U29	G4-1	G1-5	2 meters
U27	G4-0	G1-6	2 meters
U25	G3-4	G2-2	2 meters
U23	G2-5	G3-1	2 meters
U22	G0-5	G5-1	2 meters
U21	G1-2	G4-4	2 meters
U20	G0-4	G5-2	2 meters
U18	G2-4	G3-2	2 meters
U17	G1-1	G4-5	2 meters
U16	G0-3	G5-3	2 meters
U14	G2-3	G3-3	2 meters
U12	G2-2	G3-4	2 meters
U10	G2-1	G3-5	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-2	G5-4	2 meters
U4	G0-1	G5-5	2 meters
U2	G0-0	G5-6	2 meters
U1	G2-0	G3-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X7 and Earlier

Table D-4 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-4 Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U42	G5-5	G0-1	2 meters
U41	G4-5	G1-1	2 meters
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-1	G0-5	2 meters
U33	G5-0	G0-6	2 meters

Table D-4 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U31	G4-4	G1-2	2 meters
U29	G4-3	G1-3	2 meters
U27	G4-2	G1-4	2 meters
U25	G3-5	G2-1	2 meters
U23	G3-3	G2-3	2 meters
U22	G3-1	G2-5	2 meters
U21	G3-0	G2-6	2 meters
U20	G1-5	G4-1	2 meters
U18	G2-2	G3-4	2 meters
U17	G2-1	G3-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections X8

Table D-3shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-5 Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G5-5	G2-1	2 meters
U39	G3-0	G0-6	2 meters
U37	G5-6	G2-0	2 meters
U35	G4-5	G1-1	2 meters
U33	G4-4	G1-2	2 meters
U31	G3-5	G0-1	2 meters

Table D-5 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U29	G4-1	G1-5	2 meters
U27	G5-0	G2-6	2 meters
U25	G4-0	G1-6	2 meters
U23	G2-6	G5-0	2 meters
U22	G2-3	G5-3	2 meters
U21	G1-3	G4-3	2 meters
U20	G0-3	G3-3	2 meters
U18	G0-6	G3-0	2 meters
U17	G1-2	G4-4	2 meters
U16	G1-1	G4-5	2 meters
U14	G2-2	G5-4	2 meters
U12	G0-2	G3-4	2 meters
U10	G0-1	G3-5	2 meters
U8	G2-1	G5-5	2 meters
U6	G2-0	G5-6	2 meters
U4	G1-0	G4-6	2 meters
U2	G0-0	G3-6	2 meters
U1	G3-2	G0-4	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections

Table D-6 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-6 Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U42	G5-4	G2-0	2 meters
U41	G5-5	G2-1	2 meters
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters
U33	G5-2	G2-4	2 meters

Table D-6 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G4-1	G1-5	2 meters
U25	G3-5	G0-1	2 meters
U23	G3-3	G0-3	2 meters
U22	G3-2	G0-4	2 meters
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U17	G2-1	G5-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G1-5	G4-1	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Switch Cabling Connections

For InfiniBand Network Fabric, Table D-7 lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table D-7 RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
9B	9A	1 meter
10B	10A	1 meter
11B	11A	1 meter
8A	8A	1 meter
9A	9B	1 meter
10A	10B	1 meter

Table D-7 (Cont.) RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
11A	11B	1 meter
U20 8B	**U01** 1B	3 meter
U01 0B	**U22** 8B	3 meter

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections

For InfiniBand Network Fabric,Table D-8 lists the location, ports, and cables for the connections from the gateway leaf switches to the servers in a Recovery Appliance X8, X7, X6 or X5 full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table D-8 RDMA Network FabricSwitch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	6B	U41	Storage server	PCIe 7, P1	3 meters
U22	0A	U39	Storage server	PCIe 7, P1	3 meters
U22	0B	U37	Storage server	PCIe 7, P1	3 meters
U22	1A	U35	Storage server	PCIe 7, P1	3 meters
U22	1B	U33	Storage server	PCIe 7, P1	3 meters
U22	2A	U31	Storage server	PCIe 7, P1	3 meters
U22	2B	U29	Storage server	PCIe 7, P1	3 meters
U22	3B	U27	Storage server	PCIe 7, P1	3 meters
U22	4A	U25	Storage server	PCIe 7, P1	2 meters
U22	5A	U23	Storage server	PCIe 7, P1	2 meters
U22	13A	U18	Storage server	PCIe 7, P2	2 meters
U22	13B	U17	Compute server	PCIe 7, P2	2 meters
U22	14B	U16	Compute server	PCIe 7, P2	2 meters
U22	14A	U14	Storage server	PCIe 7, P2	3 meters
U22	15A	U12	Storage server	PCIe 7, P2	3 meters
U22	15B	U10	Storage server	PCIe 7, P2	3 meters
U22	16A	U8	Storage server	PCIe 7, P2	3 meters
U22	16B	U6	Storage server	PCIe 7, P2	3 meters
U22	17A	U4	Storage server	PCIe 7, P2	3 meters

Table D-8 (Cont.) RDMA Network FabricSwitch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	17B	U2	Storage server	PCIe 7, P2	3 meters
U20	6B	U41	Storage server	PCIe 7, P2	3 meters
U20	0A	U39	Storage server	PCIe 7, P2	3 meters
U20	0B	U37	Storage server	PCIe 7, P2	3 meters
U20	1A	U35	Storage server	PCIe 7, P2	3 meters
U20	1B	U33	Storage server	PCIe 7, P2	3 meters
U20	2A	U31	Storage server	PCIe 7, P2	3 meters
U20	2B	U29	Storage server	PCIe 7, P2	3 meters
U20	3B	U27	Storage server	PCIe 7, P2	3 meters
U20	4A	U25	Storage server	PCIe 7, P2	2 meters
U20	5A	U23	Storage server	PCIe 7, P2	2 meters
U20	13A	U18	Storage server	PCIe 7, P1	2 meters
U20	13B	U17	Compute server	PCIe 2, P1	2 meters
U20	14B	U16	Compute server	PCIe 2, P1	2 meters
U20	14A	U14	Storage server	PCIe 7, P1	2 meters
U20	15A	U12	Storage server	PCIe 7, P1	3 meters
U20	15B	U10	Storage server	PCIe 7, P1	3 meters
U20	16A	U8	Storage server	PCIe 7, P1	3 meters
U20	16B	U6	Storage server	PCIe 7, P1	3 meters
U20	17A	U4	Storage server	PCIe 7, P1	3 meters
U20	17B	U2	Storage server	PCIe 7, P1	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Recovery Appliance X9M Rack Layout

Figure D-3shows the front and rear views of a Recovery Appliance X9M full rack. The front is shown on the left, and the rear is shown on the right. The red bullet next to a component means that it is required for all Recovery Appliances.

Figure D-3 Rack Layout of the Recovery Appliance X9M Full Rack

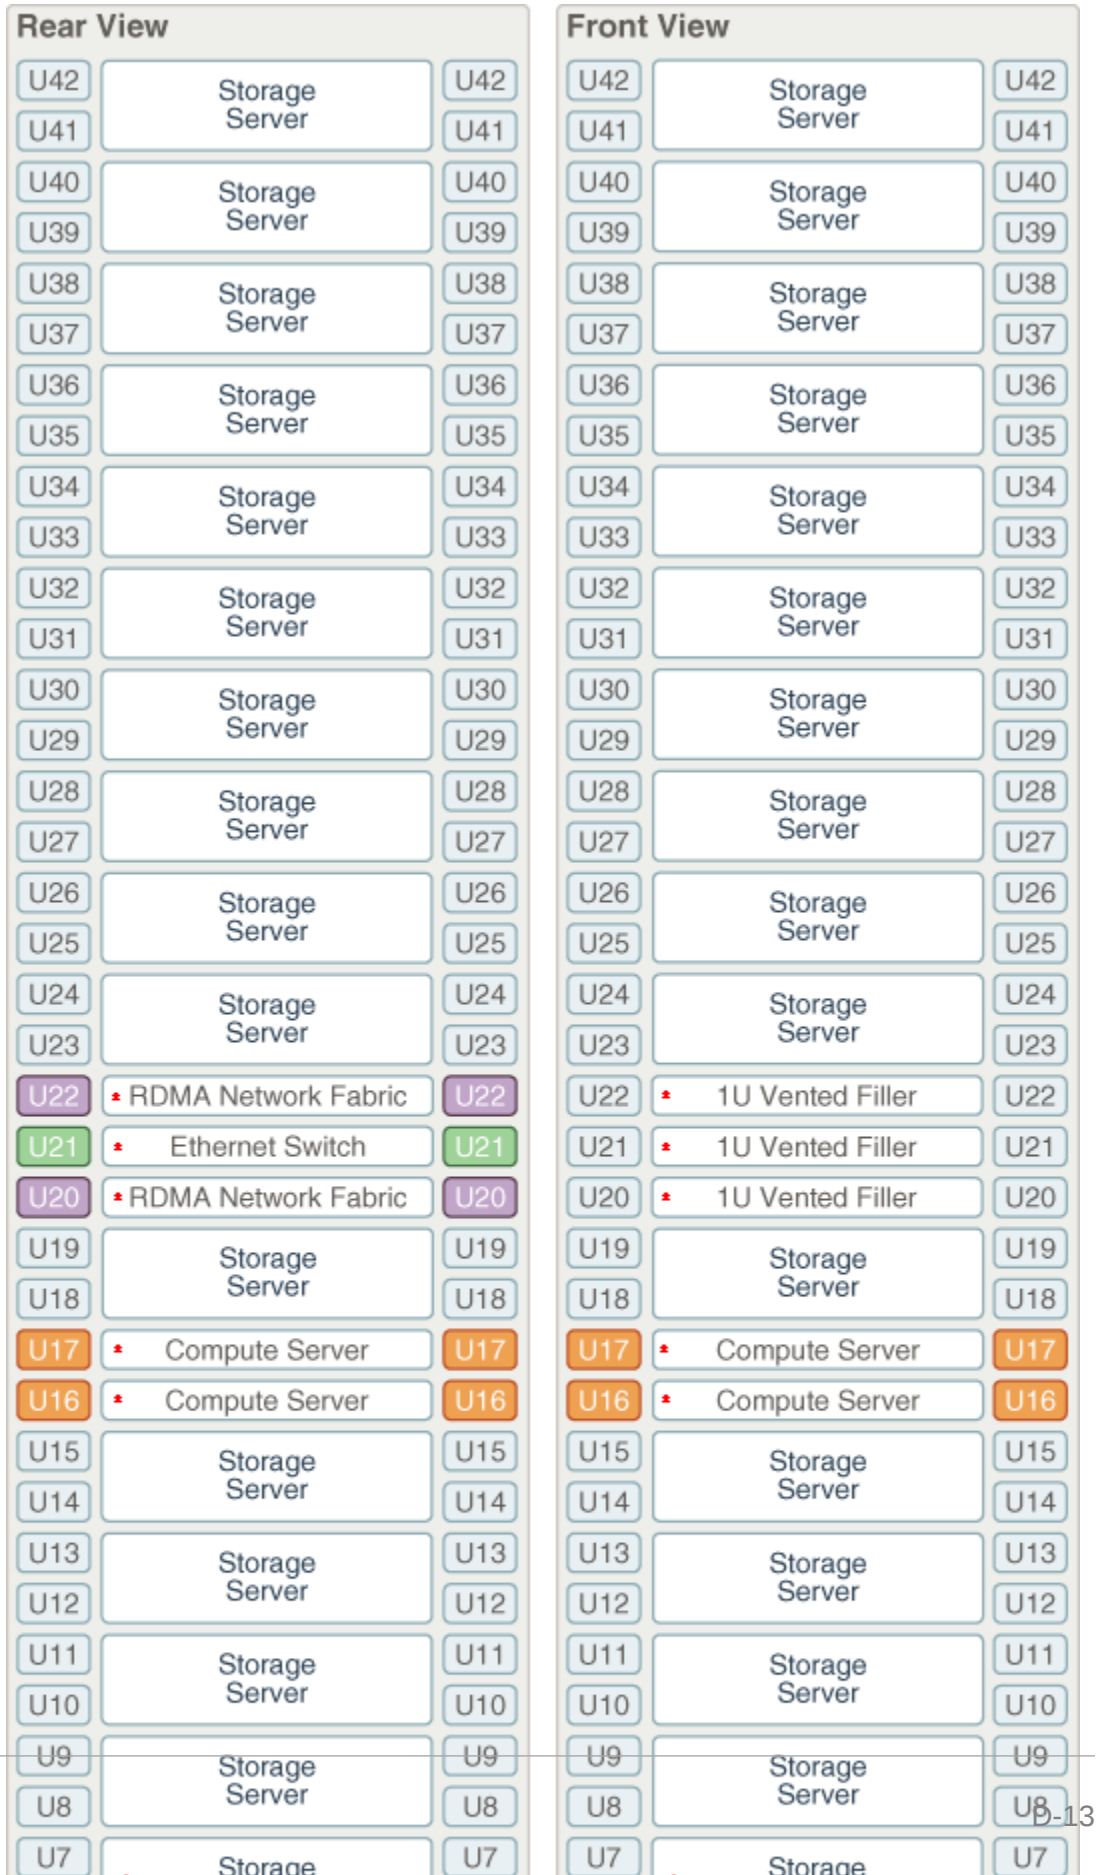

> **Note:**
>
> The 1U empty space at the bottom of the rack is reserved for an optional spine switch, which is used in multi-rack installations. Do not install any other unit in this space.

Figure E-2 shows the front and rear views of a Recovery Appliance X9M minimum configuration.

Figure D-4 Rack Layout of the Recovery Appliance X9M Minimum Configuration

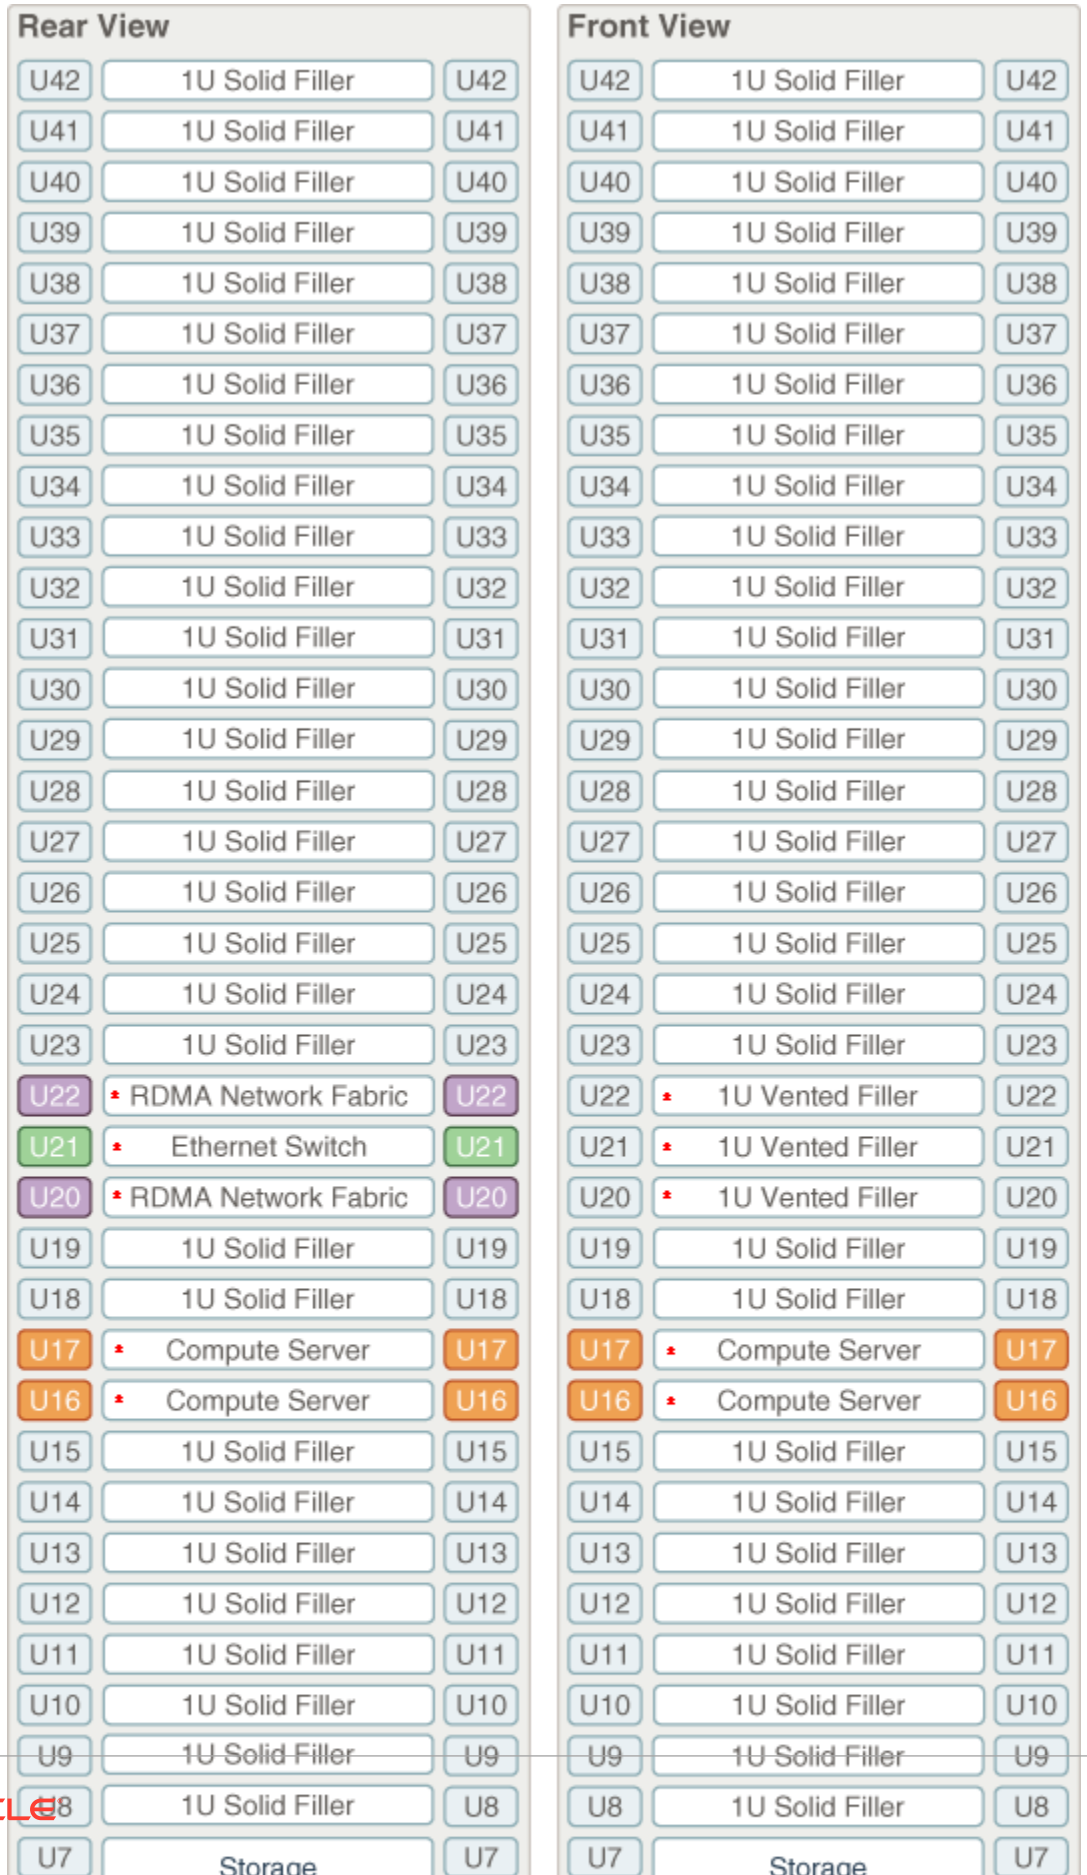

Oracle Integrated Lights Out Manager Cabling Connections X9M

Table D-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X9M full rack. The port numbers are the same in racks with fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table D-9 ILOM Cabling for the Recovery Appliance X9M Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	16
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U25	Storage server	18
U23	Storage server	22
U18	Storage server	26
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

1 U*n* is the unit location in the rack, where *n* is the number.

Administrative Ethernet Port Cabling Tables X9M

Figure D-3 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X9M full rack. The port numbers are the same in racks with fewer storage servers.

The Cisco Management Switch is located in rack unit 21 on Recovery Appliance X9M. Ethernet cables are blue.

In the following tables, U*n* denotes the unit height in the rack, where *n* is the number, such as U20. The cables used for the Cisco Management Switch are blue in color. Cabling connections marked with * are required for all rack configurations, and represent the minimum rack configuration.

Table D-10 Ethernet Cabling for the Recovery Appliance X9M Racks

From Rack Unit[1]	Network Management	To U21 Ethernet Port	Component
U39	NET 0	1	Storage server
U39	NET MGT	2	Storage server
U37	NET 0	3	Storage server
U37	NET MGT	4	Storage server
U35	NET 0	5	Storage server
U35	NET MGT	6	Storage server
U33	NET 0	7	Storage server
U33	NET MGT	8	Storage server
U31	NET 0	9	Storage server
U31	NET MGT	10	Storage server
U29	NET 0	11	Storage server
U29	NET MGT	12	Storage server
U27	NET 0	13	Storage server
U27	NET MGT	14	Storage server
U41	NET 0	15	Storage server
U41	NET MGT	16	Storage server
U25	NET 0	17	Storage server
U25	NET MGT	18	Storage server
PDU A*	NET MGT	19	
PDU B*	NET MGT	20	
U23	NET 0	21	Storage server
U23	NET MGT	22	Storage server
U18	NET 0	25	Storage server
U18	NET MGT	26	Storage server
U17*	NET 0	27	Compute server
U17*	NET MGT	28	Compute server
U16*	NET 0	29	Compute server
U16*	NET MGT	30	Compute server
U14	NET 0	31	Storage server
U14	NET MGT	32	Storage server
U12	NET 0	33	Storage server

Table D-10 (Cont.) Ethernet Cabling for the Recovery Appliance X9M Racks

From Rack Unit[1]	Network Management	To U21 Ethernet Port	Component
U12	NET MGT	34	Storage server
U10	NET 0	35	Storage server
U10	NET MGT	36	Storage server
U08	NET 0	37	Storage server
U08	NET MGT	38	Storage server
U06*	NET 0	39	Storage server
U06*	NET MGT	40	Storage server
U04*	NET 0	41	Storage server
U04*	NET MGT	42	Storage server
U02*	NET 0	43	Storage server
U02*	NET MGT	44	Storage server
U22*	NET MGT 0	45	RDMA Network Fabric switch
U20*	NET MGT 0	46	RDMA Network Fabric switch
U01*	NET MGT 0	47	
	External	48	

1 U*n* is the unit location in the rack, where *n* is the number.

* These marked components are required for all rack configurations, and together represent the minimum rack configuration.

Single-Phase Power Distribution Unit Cabling Connections X9M

Table D-11 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X9M full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-11 Single-Phase PDU Cabling for the Recovery Appliance X9M Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G4-6	G1-0	3 meters
U39	G4-4	G1-2	3 meters
U37	G5-1	G0-5	3 meters
U35	G1-5	G4-1	3 meters

Table D-11 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X9M Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U33	G2-6	G3-0	2 meters
U31	G3-3	G2-3	3 meters
U29	G3-6	G2-0	3 meters
U27	G3-2	G2-4	3 meters
U25	G4-1	G1-5	3 meters
U23	G5-0	G0-6	3meters
U22*	G1-6	G4-0	2 meters
U21*	G0-4	G5-2	2 meters
U20*	G0-3	G5-3	2 meters
U18	G4-2	G1-4	3 meters
U17*	G1-2	G4-4	3 meters
U16*	G1-4	G4-2	2 meters
U14	G2-1	G3-5	3 meters
U12	G2-0	G3-6	3 meters
U10	G2-3	G3-3	3 meters
U8	G2-4	G3-2	3 meters
U6*	G1-3	G4-3	3 meters
U4*	G0-2	G5-4	3 meters
U2*	G0-5	G5-1	3 meters
U1*	G2-5	G3-1	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

* These items are required for all Recovery Appliances.

Three-Phase Power Distribution Unit Cabling Connections X9M

Table D-12 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X9M full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-12 Three-Phase PDU Cabling for the Recovery Appliance X9M Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G5-2	G2-4	3 meters
U39	G4-6	G1-0	3 meters
U37	G4-1	G1-5	3 meters

Table D-12 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X9M Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U35	G0-4	G3-2	3 meters
U33	G2-2	G5-4	3 meters
U31	G1-2	G4-4	3 meters
U29	G0-1	G3-5	3 meters
U27	G2-0	G5-6	3 meters
U25	G3-0	G0-6	3 meters
U23	G3-5	G0-1	3 meters
U22*	G1-6	G4-0	2 meters
U21*	G2-6	G5-0	2 meters
U20*	G1-5	G4-1	2 meters
U18	G1-4	G4-2	3 meters
U17*	G2-1	G5-5	3 meters
U16*	G1-3	G4-3	3 meters
U14	G0-5	G3-1	3 meters
U12	G2-5	G5-1	3 meters
U10	G0-3	G3-3	3 meters
U8	G2-3	G5-3	3 meters
U6*	G2-4	G5-2	3 meters
U4*	G0-6	G3-0	3 meters
U2*	G0-0	G3-6	3 meters
U1*	G1-0	G4-6	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

* These items are required for all Recovery Appliances.

RDMA Network Fabric Switch-to-Switch Cabling Connections X9M

For RoCE Network Fabric, Table D-7 lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table D-13 RDMA Network Fabric Switch-to-Switch Cabling X9M

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
04	04	1 meter
05	05	1 meter
06	06	1 meter
07	07	1 meter
30	30	1 meter
31	31	1 meter
32	32	1 meter
33	33	3 meter

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections X9M

The following tables list the location, ports, and cables for the connections from the RDMA Network Fabric switches to the computer and HC storage servers in a Recovery Appliance X9M full rack. The port numbers and cable lengths are the same in racks with fewer HC Storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table D-14 RDMA Network Fabric Switch to Compute Server Cabling for the Recovery Appliance X9M Full Racks

From RDMA Network Fabric Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U20	21*	U17	1U compute server	PCIe 5, P1	3 meters
U20	22*	U16	1U compute server	PCIe 5, P1	3 meters
U22	21*	U17	1U compute server	PCIe 5, P2	3 meters
U22	22*	U16	1U compute server	PCIe 5, P2	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

* All Recovery Appliances have these entries.

Table D-15 RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance X9M Full Racks

From RDMA Network Fabric Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U20	8	U39	HC Storage server	PCIe 10, P1	3 meters
U20	9	U37	HC Storage server	PCIe 10, P1	3 meters
U20	10	U29	HC Storage server	PCIe 10, P1	3 meters
U20	11	U27	HC Storage server	PCIe 10, P1	3 meters
U20	12	U25	HC Storage server	PCIe 10, P1	2 meters
U20	13	U23	HC Storage server	PCIe 10, P1	2 meters
U20	14	U18	HC Storage server	PCIe 10, P1	2 meters
U20	15	U17	HC Storage server	PCIe 10, P1	2 meters
U20	16	U16	HC Storage server	PCIe 10, P1	2 meters
U20	18	U14	HC Storage server	PCIe 10, P1	3 meters
U20	20	U12	HC Storage server	PCIe 10, P1	3 meters
U20	23	U10	HC Storage server	PCIe 10, P1	3 meters
U20	24	U8	HC Storage server	PCIe 10, P1	3 meters
U20	25	U6	HC Storage server	PCIe 10, P1	3 meters
U20	26	U4	HC Storage server	PCIe 10, P1	3 meters
U20	27*	U2	HC Storage server	PCIe 10, P1	3 meters
U20	28*	U41	HC Storage server	PCIe 10, P1	3 meters
U20	29*	U39	HC Storage server	PCIe 10, P1	3 meters
U22	08	U39	HC Storage server	PCIe 10, P2	3 meters
U22	09	U37	HC Storage server	PCIe 10, P2	3 meters
U22	10	U35	HC Storage server	PCIe 10, P2	3 meters

Table D-15 (Cont.) RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance X9M Full Racks

From RDMA Network Fabric Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	11	U33	HC Storage server	PCIe 10, P2	3 meters
U22	12	U31	HC Storage server	PCIe 10, P2	3 meters
U22	13	U29	HC Storage server	PCIe 10, P2	3 meters
U22	14	U27	HC Storage server	PCIe 10, P2	2 meters
U22	15	U41	HC Storage server	PCIe 10, P2	2 meters
U22	16	U25	HC Storage server	PCIe 10, P2	2 meters
U22	18	U23	HC Storage server	PCIe 10, P2	2 meters
U22	20	U18	HC Storage server	PCIe 10, P2	2 meters
U22	23	U14	HC Storage server	PCIe 10, P2	2 meters
U22	24	U12	HC Storage server	PCIe 10, P2	2 meters
U22	25	U10	HC Storage server	PCIe 10, P2	2 meters
U22	26	U08	HC Storage server	PCIe 10, P2	2 meters
U22	27*	U06	HC Storage server	PCIe 10, P2	2 meters
U22	28*	U04	HC Storage server	PCIe 10, P2	2 meters
U22	29*	U02	HC Storage server	PCIe 10, P2	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

* All Recovery Appliances have these entries.

E

In-Rack Cabling Tables for Recovery Appliance X8M

This appendix identifies the location of the various components in a Recovery Appliance X8M full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black**: RDMA Network Fabric cables, AC power jumper cables, and Ethernet cables for storage servers
- **Blue**: Gigabit Ethernet cables for compute servers
- **Red**: Integrated Lights Out Manager (ILOM) cables

Recovery Appliance X9M and X8M Rack Layout

Figure E-1shows the front and rear views of a Recovery Appliance X9M and X8M full rack. The front is shown on the left, and the rear is shown on the right.

> **Note:**
>
> The 1U empty space at the bottom of the rack is reserved for an optional spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure E-1 Rack Layout of the Recovery Appliance X9M and X8M Full Rack

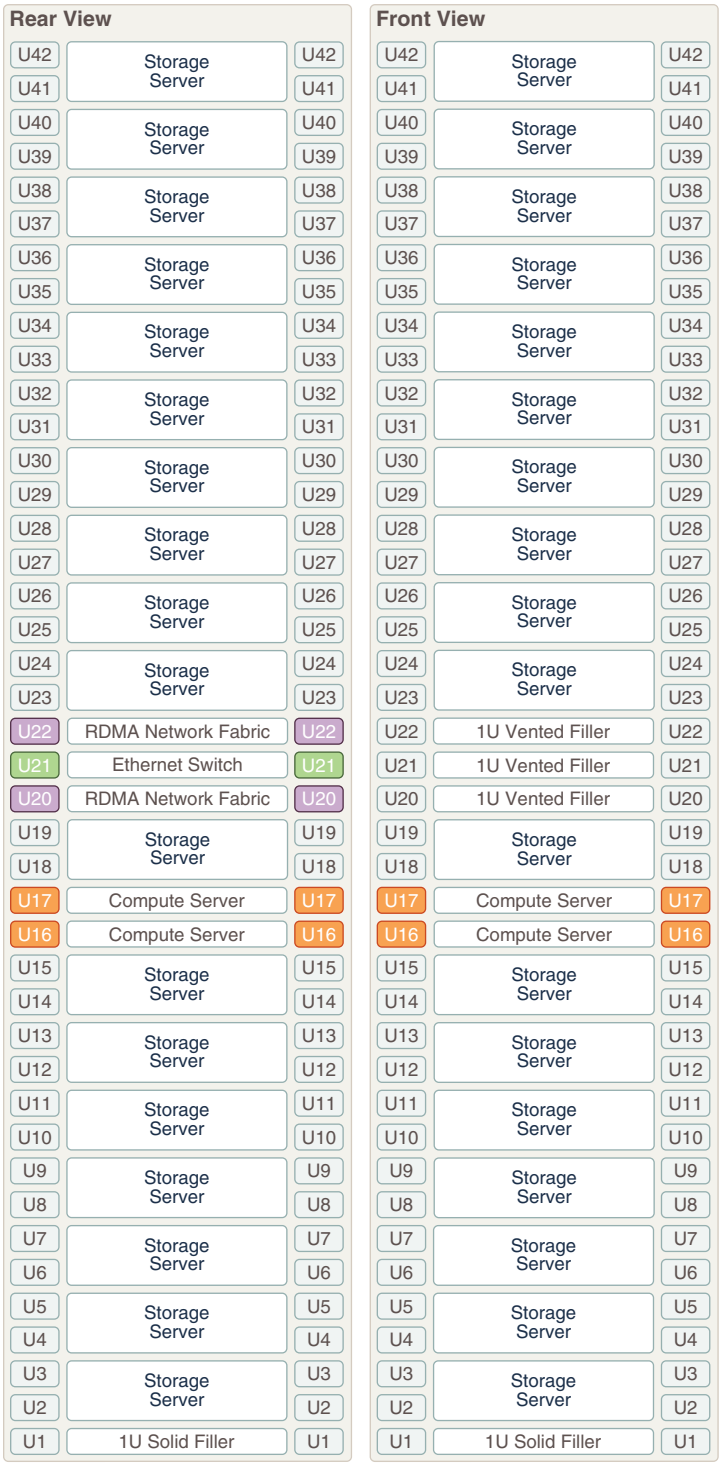

Figure E-2 shows the front and rear views of a Recovery Appliance X9M and X8M minimum configuration.

Figure E-2 Rack Layout of the Recovery Appliance X9M and X8M Minimum Configuration

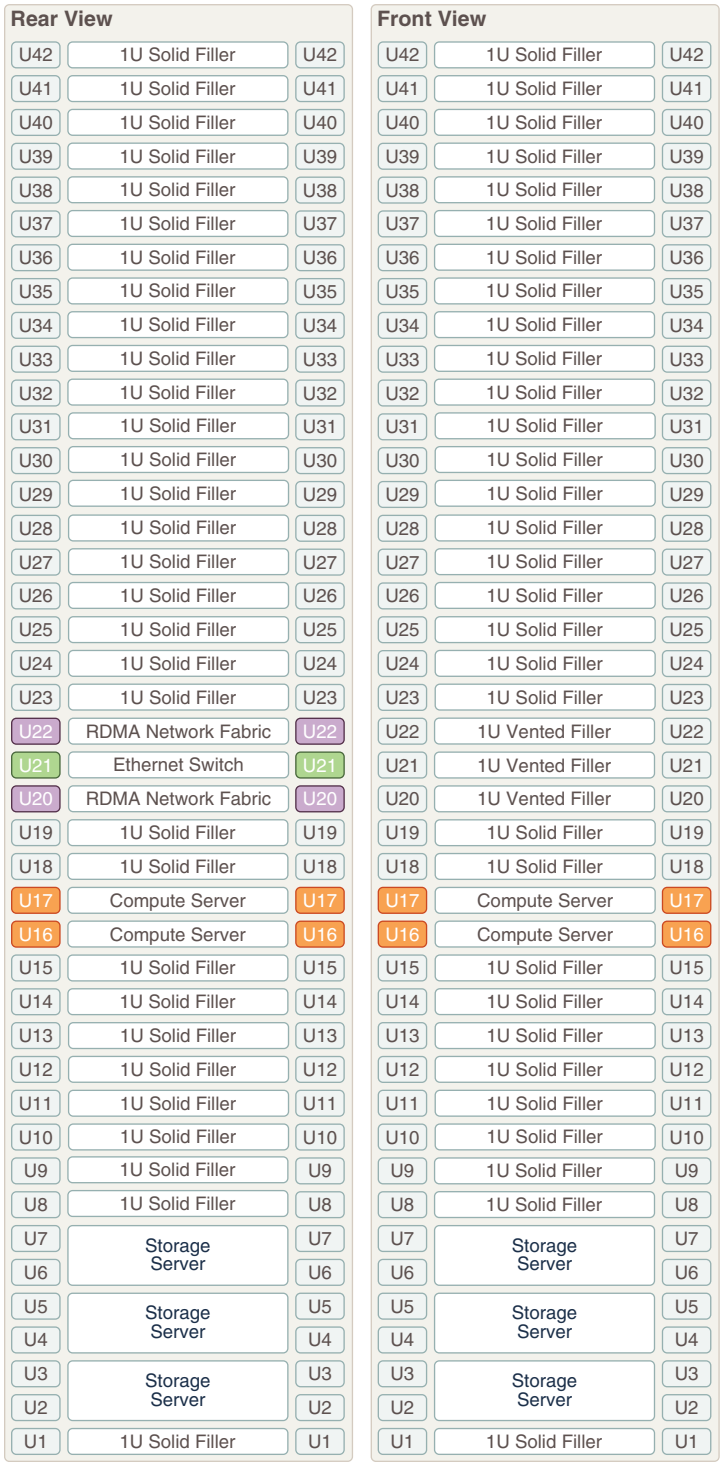

Oracle Integrated Lights Out Manager Cabling Connections X8M

Table D-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X8M full rack. The port numbers are the same in racks with fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table E-1 ILOM Cabling for the Recovery Appliance X8M Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	16
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U25	Storage server	18
U23	Storage server	22
U18	Storage server	26
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

1 U*n* is the unit location in the rack, where *n* is the number.

Administrative Ethernet Port Cabling Tables X8M

Table E-2 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X8M full rack. The port numbers are the same in racks with fewer storage servers.

The Cisco Management Switch is located in rack unit 21 on Recovery Appliance X8M. Ethernet cables are blue.

In the following tables, U*n* denotes the unit height in the rack, where *n* is the number, such as U20. The cables used for the Cisco Management Switch are blue in color. The following tables are in this section:

Table E-2 Ethernet Cabling for the Recovery Appliance X8M Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	15
U39	Storage server	1
U37	Storage server	3
U35	Storage server	5
U33	Storage server	7
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U25	Storage server	17
U23	Storage server	21
U22	RDMA Network Fabric switch	45
U20	RDMA Network Fabric switch	46
U18	Storage server	25
U17	Compute server	27
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X8M

Table D-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8M full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table E-3 Single-Phase PDU Cabling for the Recovery Appliance X8M Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G4-2	G1-4	2 meters
U39	G4-4	G1-2	2 meters
U37	G5-3	G2-3	2 meters
U35	G3-2	G0-4	2 meters
U33	G5-4	G2-2	2 meters
U31	G5-1	G2-5	2 meters
U29	G3-0	G0-6	2 meters
U27	G3-5	G0-1	2 meters
U25	G4-5	G1-1	2 meters
U23	G4-0	G1-6	2 meters
U22	G1-6	G4-0	2 meters
U21	G2-6	G5-0	2 meters
U20	G1-5	G5-1	2 meters
U18	G1-4	G4-2	2 meters
U17	G2-3	G4-3	2 meters
U16	G2-0	G5-6	2 meters
U14	G1-0	G4-6	2 meters
U12	G0-2	G3-4	2 meters
U10	G0-1	G3-5	2 meters
U8	G0-5	G3-1	2 meters
U6	G2-4	G5-2	2 meters
U4	G0-3	G3-3	2 meters
U2	G1-3	G4-3	2 meters
U1	G1-1	G4-5	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections X8M

Table D-3shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8M full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table E-4 Three-Phase PDU Cabling for the Recovery Appliance X8M Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G4-4	G1-2	2 meters
U39	G5-1	G0-5	2 meters
U37	G5-5	G0-1	2 meters
U35	G4-1	G1-5	2 meters
U33	G5-3	G0-3	2 meters
U31	G5-4	G0-2	2 meters
U29	G4-6	G1-0	2 meters
U27	G3-4	G2-2	2 meters
U25	G3-1	G2-5	2 meters
U23	G3-0	G2-6	2 meters
U22	G1-6	G4-0	2 meters
U21	G1-4	G4-2	2 meters
U20	G1-5	G4-1	2 meters
U18	G2-6	G3-0	2 meters
U17	G0-5	G5-1	2 meters
U16	G1-1	G4-5	2 meters
U14	G2-2	G3-4	2 meters
U12	G2-5	G3-1	2 meters
U10	G2-4	G3-2	2 meters
U8	G0-2	G5-4	2 meters
U6	G1-0	G4-6	2 meters
U4	G0-0	G5-6	2 meters
U2	G0-4	G5-2	2 meters
U1	G4-2	G1-4	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Switch Cabling ConnectionsX8M

For RoCE Network Fabric, Table D-7 lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table E-5 RDMA Network Fabric Switch-to-Switch Cabling X8M

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
04	04	1 meter
05	05	1 meter
06	06	1 meter
07	07	1 meter
30	30	1 meter
31	31	1 meter
32	32	1 meter
33	33	3 meter

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections X8M

Table D-8 lists the location, ports, and cables for the InfiniBand connections from the gateway leaf switches to the servers in a Recovery Appliance X8M full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table E-6 RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8M Full Racks

From RDMA Network Fabric Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	15	U41	Storage server	PCIe 1, P2	3 meters
U22	08	U39	Storage server	PCIe 1, P2	3 meters
U22	09	U37	Storage server	PCIe 1, P2	3 meters
U22	10	U35	Storage server	PCIe 1, P2	3 meters
U22	11	U33	Storage server	PCIe 1, P2	3 meters
U22	12	U31	Storage server	PCIe 1, P2	3 meters
U22	13	U29	Storage server	PCIe 1, P2	3 meters
U22	14	U27	Storage server	PCIe 1, P2	3 meters
U22	16	U25	Storage server	PCIe 1, P2	2 meters
U22	18	U23	Storage server	PCIe 1, P2	2 meters
U22	20	U18	Storage server	PCIe 1, P2	2 meters
U22	21	U17	Compute server	PCIe 1, P2	2 meters

Table E-6 (Cont.) RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8M Full Racks

From RDMA Network Fabric Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	22	U16	Compute server	PCIe 1, P2	2 meters
U22	23	U14	Storage server	PCIe 1, P2	3 meters
U22	24	U12	Storage server	PCIe 1, P2	3 meters
U22	25	U10	Storage server	PCIe 1, P2	3 meters
U22	26	U8	Storage server	PCIe 1, P2	3 meters
U22	27	U6	Storage server	PCIe 1, P2	3 meters
U22	28	U4	Storage server	PCIe 1, P2	3 meters
U22	29	U2	Storage server	PCIe 1, P2	3 meters
U20	15	U41	Storage server	PCIe 1, P1	3 meters
U20	08	U39	Storage server	PCIe 1, P1	3 meters
U20	09	U37	Storage server	PCIe 1, P1	3 meters
U20	10	U35	Storage server	PCIe 1, P1	3 meters
U20	11	U33	Storage server	PCIe 1, P1	3 meters
U20	12	U31	Storage server	PCIe 1, P1	3 meters
U20	13	U29	Storage server	PCIe 1, P1	3 meters
U20	14	U27	Storage server	PCIe 1, P1	3 meters
U20	16	U25	Storage server	PCIe 1, P1	2 meters
U20	18	U23	Storage server	PCIe 1, P1	2 meters
U20	20	U18	Storage server	PCIe 1, P1	2 meters
U20	21	U17	Compute server	PCIe 1, P1	2 meters
U20	22	U16	Compute server	PCIe 1, P1	2 meters
U20	23	U14	Storage server	PCIe 1, P1	2 meters
U20	24	U12	Storage server	PCIe 1, P1	3 meters
U20	25	U10	Storage server	PCIe 1, P1	3 meters
U20	26	U8	Storage server	PCIe 1, P1	3 meters
U20	27	U6	Storage server	PCIe 1, P1	3 meters
U20	28	U4	Storage server	PCIe 1, P1	3 meters
U20	29	U2	Storage server	PCIe 1, P1	3 meters

[1] U*n* is the unit location in the rack, where *n* is the number.

F

In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5

This appendix identifies the location of the various components in a Recovery Appliance X8, X7, X6, and X5 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

This appendix contains the following sections:

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black**: RDMA Network Fabric cables, AC power jumper cables, and Ethernet cables for storage servers
- **Blue**: Gigabit Ethernet cables for compute servers
- **Red**: Integrated Lights Out Manager (ILOM) cables

Recovery Appliance X8, X7, X6, and X5 Rack Layout

Figure D-1 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 full rack. The front is shown on the left, and the rear is shown on the right.

> **Note:**
>
> The 1U empty space at the bottom of the rack is reserved for an optional InfiniBand spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure F-1 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Full Rack

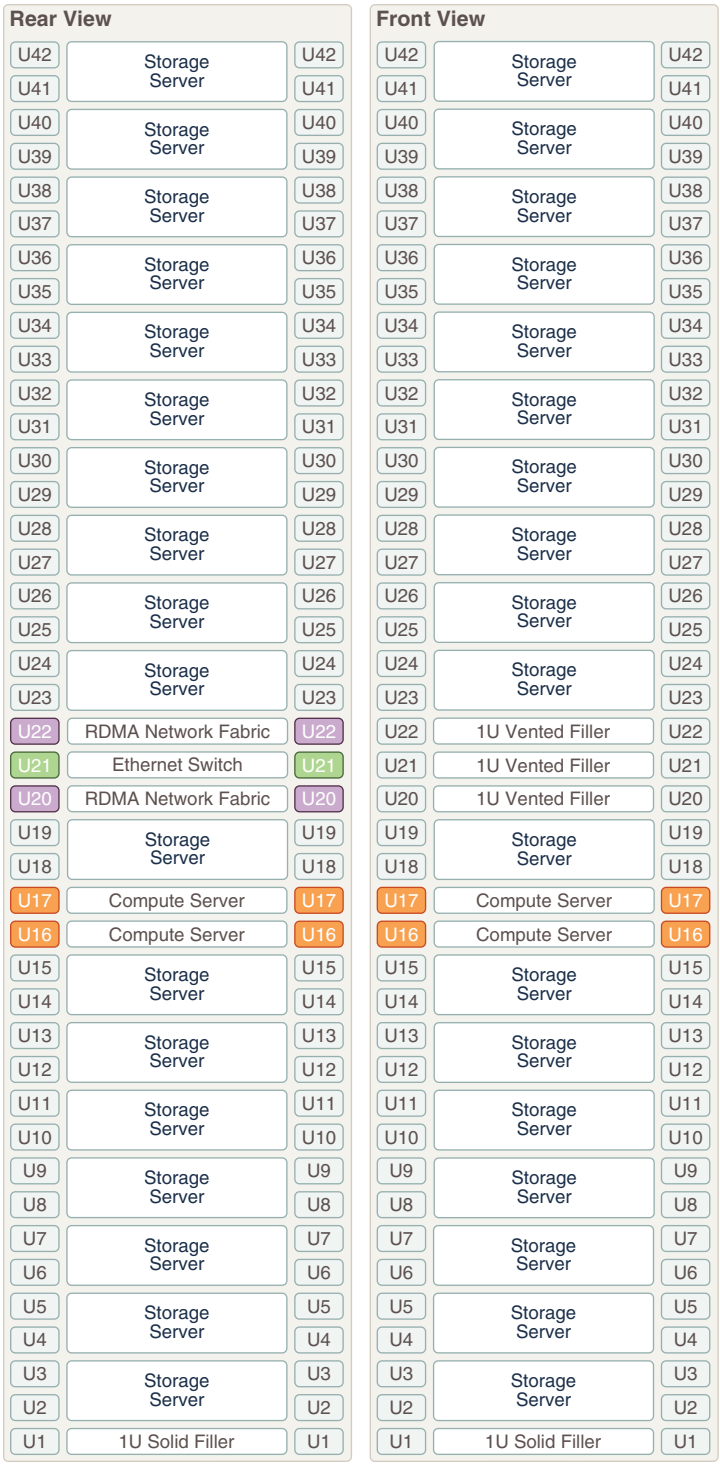

Figure D-2 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 minimum configuration.

Figure F-2 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Minimum Configuration

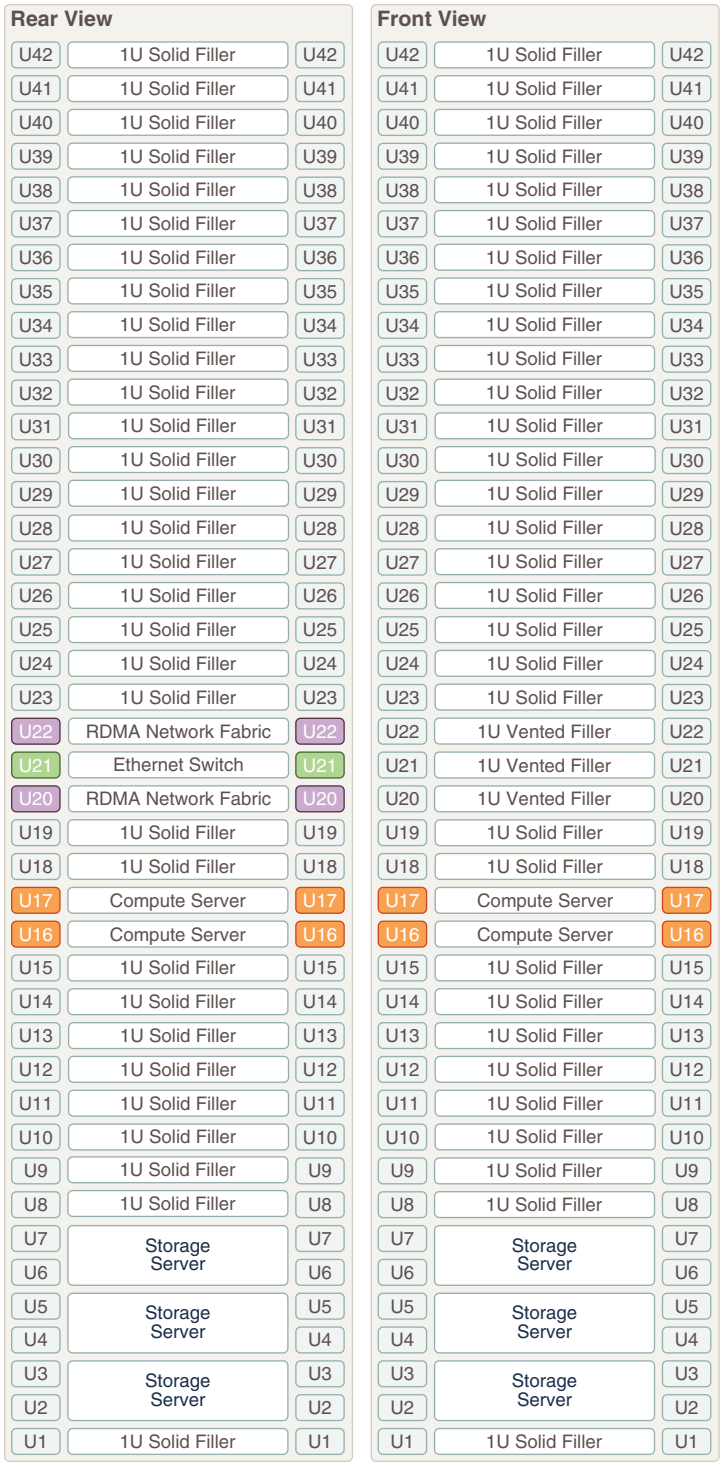

Oracle Integrated Lights Out Manager Cabling Connections

Table D-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in racks with

fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table F-1 ILOM Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	16
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U25	Storage server	18
U23	Storage server	22
U18	Storage server	26
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

1 U*n* is the unit location in the rack, where *n* is the number.

Administrative Gigabit Ethernet Cabling Connections

Table D-2 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in racks with fewer storage servers. The Ethernet switch is located in rack unit 21. Ethernet cables are blue.

Table F-2 Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	15
U39	Storage server	1
U37	Storage server	3

Table F-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U35	Storage server	5
U33	Storage server	7
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U25	Storage server	17
U23	Storage server	21
U22	InfiniBand leaf switch	45
U20	InfiniBand leaf switch	46
U18	Storage server	25
U17	Compute server	27
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X8

Table D-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-3 Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G3-6	G2-0	2 meters
U39	G5-4	G0-2	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-2	G0-4	2 meters
U33	G5-0	G0-6	2 meters
U31	G4-2	G1-4	2 meters

Table F-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U29	G4-1	G1-5	2 meters
U27	G4-0	G1-6	2 meters
U25	G3-4	G2-2	2 meters
U23	G2-5	G3-1	2 meters
U22	G0-5	G5-1	2 meters
U21	G1-2	G4-4	2 meters
U20	G0-4	G5-2	2 meters
U18	G2-4	G3-2	2 meters
U17	G1-1	G4-5	2 meters
U16	G0-3	G5-3	2 meters
U14	G2-3	G3-3	2 meters
U12	G2-2	G3-4	2 meters
U10	G2-1	G3-5	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-2	G5-4	2 meters
U4	G0-1	G5-5	2 meters
U2	G0-0	G5-6	2 meters
U1	G2-0	G3-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X7 and Earlier

Table D-4 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-4 Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U42	G5-5	G0-1	2 meters
U41	G4-5	G1-1	2 meters
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters

Table F-4 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U35	G5-1	G0-5	2 meters
U33	G5-0	G0-6	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-3	G1-3	2 meters
U27	G4-2	G1-4	2 meters
U25	G3-5	G2-1	2 meters
U23	G3-3	G2-3	2 meters
U22	G3-1	G2-5	2 meters
U21	G3-0	G2-6	2 meters
U20	G1-5	G4-1	2 meters
U18	G2-2	G3-4	2 meters
U17	G2-1	G3-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections X8

Table D-3shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-5 Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G5-5	G2-1	2 meters
U39	G3-0	G0-6	2 meters
U37	G5-6	G2-0	2 meters
U35	G4-5	G1-1	2 meters

Table F-5 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U33	G4-4	G1-2	2 meters
U31	G3-5	G0-1	2 meters
U29	G4-1	G1-5	2 meters
U27	G5-0	G2-6	2 meters
U25	G4-0	G1-6	2 meters
U23	G2-6	G5-0	2 meters
U22	G2-3	G5-3	2 meters
U21	G1-3	G4-3	2 meters
U20	G0-3	G3-3	2 meters
U18	G0-6	G3-0	2 meters
U17	G1-2	G4-4	2 meters
U16	G1-1	G4-5	2 meters
U14	G2-2	G5-4	2 meters
U12	G0-2	G3-4	2 meters
U10	G0-1	G3-5	2 meters
U8	G2-1	G5-5	2 meters
U6	G2-0	G5-6	2 meters
U4	G1-0	G4-6	2 meters
U2	G0-0	G3-6	2 meters
U1	G3-2	G0-4	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections

Table D-6 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-6 Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U42	G5-4	G2-0	2 meters
U41	G5-5	G2-1	2 meters
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters

Table F-6 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U33	G5-2	G2-4	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G4-1	G1-5	2 meters
U25	G3-5	G0-1	2 meters
U23	G3-3	G0-3	2 meters
U22	G3-2	G0-4	2 meters
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U17	G2-1	G5-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G1-5	G4-1	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Switch Cabling Connections

For InfiniBand Network Fabric, Table D-7 lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table F-7 RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
9B	9A	1 meter
10B	10A	1 meter
11B	11A	1 meter
8A	8A	1 meter
9A	9B	1 meter

Table F-7 (Cont.) RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
10A	10B	1 meter
11A	11B	1 meter
U20 8B	**U01** 1B	3 meter
U01 0B	**U22** 8B	3 meter

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections

For InfiniBand Network Fabric,Table D-8 lists the location, ports, and cables for the connections from the gateway leaf switches to the servers in a Recovery Appliance X8, X7, X6 or X5 full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table F-8 RDMA Network FabricSwitch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	6B	U41	Storage server	PCIe 7, P1	3 meters
U22	0A	U39	Storage server	PCIe 7, P1	3 meters
U22	0B	U37	Storage server	PCIe 7, P1	3 meters
U22	1A	U35	Storage server	PCIe 7, P1	3 meters
U22	1B	U33	Storage server	PCIe 7, P1	3 meters
U22	2A	U31	Storage server	PCIe 7, P1	3 meters
U22	2B	U29	Storage server	PCIe 7, P1	3 meters
U22	3B	U27	Storage server	PCIe 7, P1	3 meters
U22	4A	U25	Storage server	PCIe 7, P1	2 meters
U22	5A	U23	Storage server	PCIe 7, P1	2 meters
U22	13A	U18	Storage server	PCIe 7, P2	2 meters
U22	13B	U17	Compute server	PCIe 7, P2	2 meters
U22	14B	U16	Compute server	PCIe 7, P2	2 meters
U22	14A	U14	Storage server	PCIe 7, P2	3 meters
U22	15A	U12	Storage server	PCIe 7, P2	3 meters
U22	15B	U10	Storage server	PCIe 7, P2	3 meters
U22	16A	U8	Storage server	PCIe 7, P2	3 meters

Table F-8 (Cont.) RDMA Network FabricSwitch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	16B	U6	Storage server	PCIe 7, P2	3 meters
U22	17A	U4	Storage server	PCIe 7, P2	3 meters
U22	17B	U2	Storage server	PCIe 7, P2	3 meters
U20	6B	U41	Storage server	PCIe 7, P2	3 meters
U20	0A	U39	Storage server	PCIe 7, P2	3 meters
U20	0B	U37	Storage server	PCIe 7, P2	3 meters
U20	1A	U35	Storage server	PCIe 7, P2	3 meters
U20	1B	U33	Storage server	PCIe 7, P2	3 meters
U20	2A	U31	Storage server	PCIe 7, P2	3 meters
U20	2B	U29	Storage server	PCIe 7, P2	3 meters
U20	3B	U27	Storage server	PCIe 7, P2	3 meters
U20	4A	U25	Storage server	PCIe 7, P2	2 meters
U20	5A	U23	Storage server	PCIe 7, P2	2 meters
U20	13A	U18	Storage server	PCIe 7, P1	2 meters
U20	13B	U17	Compute server	PCIe 2, P1	2 meters
U20	14B	U16	Compute server	PCIe 2, P1	2 meters
U20	14A	U14	Storage server	PCIe 7, P1	2 meters
U20	15A	U12	Storage server	PCIe 7, P1	3 meters
U20	15B	U10	Storage server	PCIe 7, P1	3 meters
U20	16A	U8	Storage server	PCIe 7, P1	3 meters
U20	16B	U6	Storage server	PCIe 7, P1	3 meters
U20	17A	U4	Storage server	PCIe 7, P1	3 meters
U20	17B	U2	Storage server	PCIe 7, P1	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5

This appendix identifies the location of the various components in a Recovery Appliance X8, X7, X6, and X5 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

This appendix contains the following sections:

- Cable Color Coding
- Recovery Appliance X8, X7, X6, and X5 Rack Layout
- Oracle Integrated Lights Out Manager Cabling Connections

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black**: RDMA Network Fabric cables, AC power jumper cables, and Ethernet cables for storage servers
- **Blue**: Gigabit Ethernet cables for compute servers
- **Red**: Integrated Lights Out Manager (ILOM) cables

Recovery Appliance X8, X7, X6, and X5 Rack Layout

Figure D-1 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 full rack. The front is shown on the left, and the rear is shown on the right.

> **Note:**
>
> The 1U empty space at the bottom of the rack is reserved for an optional InfiniBand spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure F-3 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Full Rack

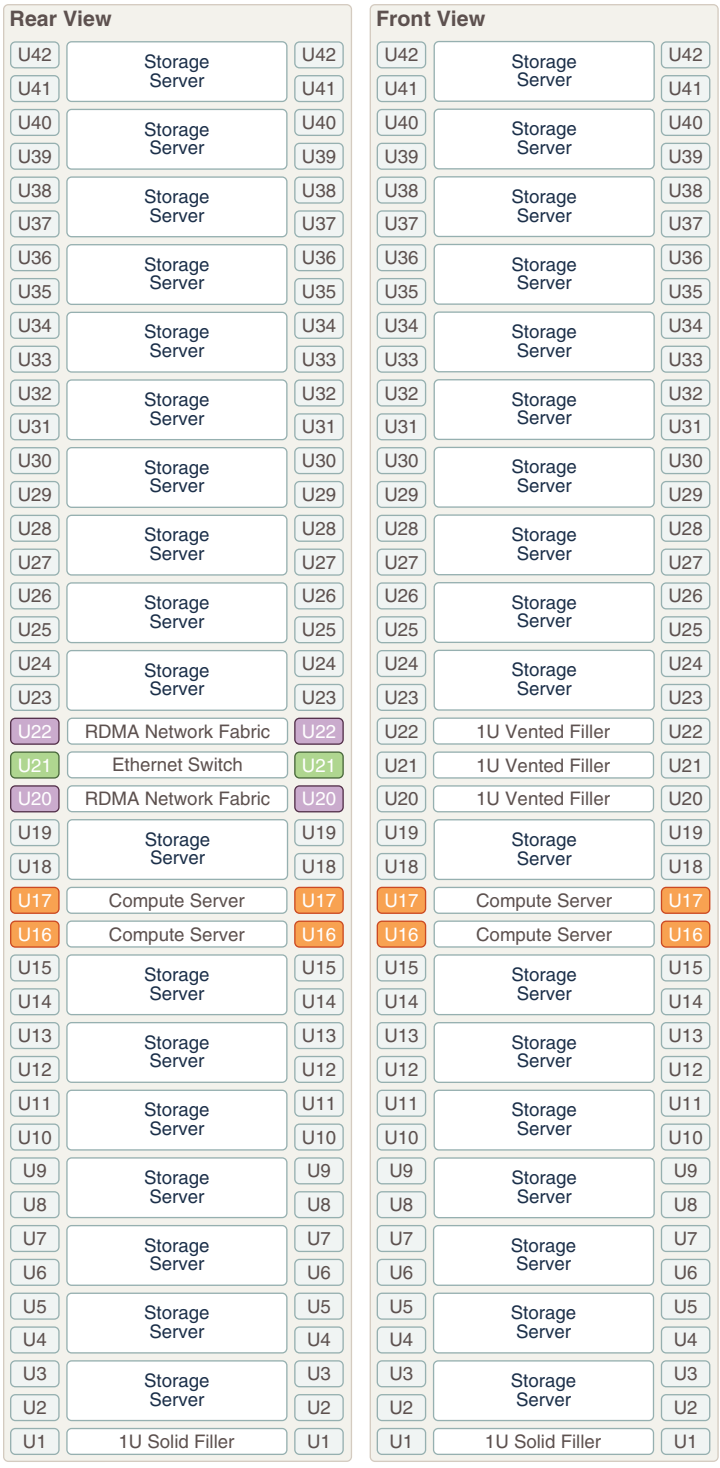

Figure D-2 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 minimum configuration.

Figure F-4 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Minimum Configuration

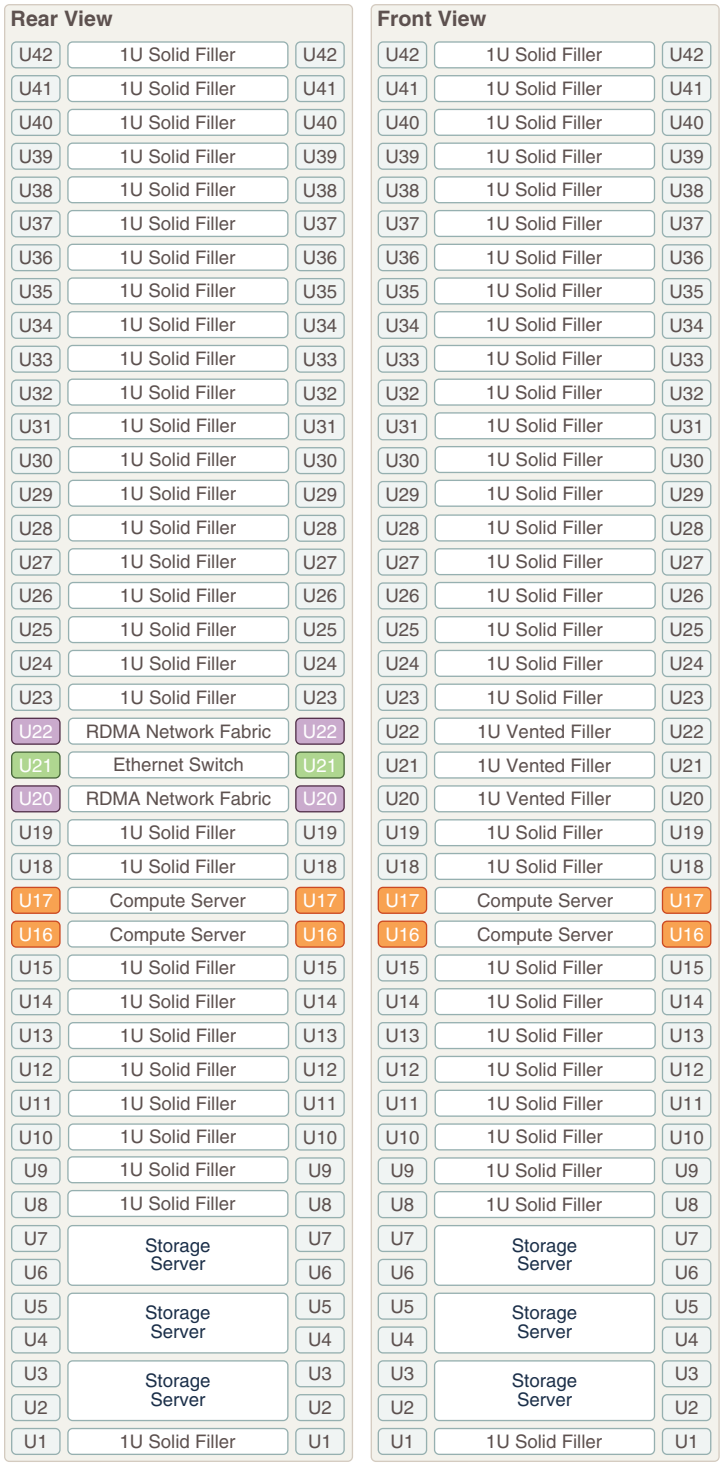

Oracle Integrated Lights Out Manager Cabling Connections

Table D-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in

racks with fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table F-9 ILOM Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	16
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U25	Storage server	18
U23	Storage server	22
U18	Storage server	26
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

1 U*n* is the unit location in the rack, where *n* is the number.

Administrative Gigabit Ethernet Cabling Connections

Table D-2 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in racks with fewer storage servers. The Ethernet switch is located in rack unit 21. Ethernet cables are blue.

Table F-10 Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U41	Storage server	15
U39	Storage server	1
U37	Storage server	3

Table F-10 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U35	Storage server	5
U33	Storage server	7
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U25	Storage server	17
U23	Storage server	21
U22	InfiniBand leaf switch	45
U20	InfiniBand leaf switch	46
U18	Storage server	25
U17	Compute server	27
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X8

Table D-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-11 Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G3-6	G2-0	2 meters
U39	G5-4	G0-2	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-2	G0-4	2 meters
U33	G5-0	G0-6	2 meters

Table F-11 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U31	G4-2	G1-4	2 meters
U29	G4-1	G1-5	2 meters
U27	G4-0	G1-6	2 meters
U25	G3-4	G2-2	2 meters
U23	G2-5	G3-1	2 meters
U22	G0-5	G5-1	2 meters
U21	G1-2	G4-4	2 meters
U20	G0-4	G5-2	2 meters
U18	G2-4	G3-2	2 meters
U17	G1-1	G4-5	2 meters
U16	G0-3	G5-3	2 meters
U14	G2-3	G3-3	2 meters
U12	G2-2	G3-4	2 meters
U10	G2-1	G3-5	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-2	G5-4	2 meters
U4	G0-1	G5-5	2 meters
U2	G0-0	G5-6	2 meters
U1	G2-0	G3-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X7 and Earlier

Table D-4 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-12 Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U42	G5-5	G0-1	2 meters
U41	G4-5	G1-1	2 meters
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-1	G0-5	2 meters

Table F-12 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U33	G5-0	G0-6	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-3	G1-3	2 meters
U27	G4-2	G1-4	2 meters
U25	G3-5	G2-1	2 meters
U23	G3-3	G2-3	2 meters
U22	G3-1	G2-5	2 meters
U21	G3-0	G2-6	2 meters
U20	G1-5	G4-1	2 meters
U18	G2-2	G3-4	2 meters
U17	G2-1	G3-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections X8

Table D-3shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-13 Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U41	G5-5	G2-1	2 meters
U39	G3-0	G0-6	2 meters
U37	G5-6	G2-0	2 meters
U35	G4-5	G1-1	2 meters
U33	G4-4	G1-2	2 meters

Table F-13 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U31	G3-5	G0-1	2 meters
U29	G4-1	G1-5	2 meters
U27	G5-0	G2-6	2 meters
U25	G4-0	G1-6	2 meters
U23	G2-6	G5-0	2 meters
U22	G2-3	G5-3	2 meters
U21	G1-3	G4-3	2 meters
U20	G0-3	G3-3	2 meters
U18	G0-6	G3-0	2 meters
U17	G1-2	G4-4	2 meters
U16	G1-1	G4-5	2 meters
U14	G2-2	G5-4	2 meters
U12	G0-2	G3-4	2 meters
U10	G0-1	G3-5	2 meters
U8	G2-1	G5-5	2 meters
U6	G2-0	G5-6	2 meters
U4	G1-0	G4-6	2 meters
U2	G0-0	G3-6	2 meters
U1	G3-2	G0-4	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections

Table D-6 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-14 Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U42	G5-4	G2-0	2 meters
U41	G5-5	G2-1	2 meters
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters

Table F-14 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U33	G5-2	G2-4	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G4-1	G1-5	2 meters
U25	G3-5	G0-1	2 meters
U23	G3-3	G0-3	2 meters
U22	G3-2	G0-4	2 meters
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U17	G2-1	G5-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G1-5	G4-1	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Switch Cabling Connections

For InfiniBand Network Fabric, Table D-7 lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table F-15 RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
9B	9A	1 meter
10B	10A	1 meter
11B	11A	1 meter
8A	8A	1 meter
9A	9B	1 meter

Table F-15 (Cont.) RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port[1]	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
10A	10B	1 meter
11A	11B	1 meter
U20 8B	**U01** 1B	3 meter
U01 0B	**U22** 8B	3 meter

1 U*n* is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections

For InfiniBand Network Fabric,Table D-8 lists the location, ports, and cables for the connections from the gateway leaf switches to the servers in a Recovery Appliance X8, X7, X6 or X5 full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table F-16 RDMA Network FabricSwitch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	6B	U41	Storage server	PCIe 7, P1	3 meters
U22	0A	U39	Storage server	PCIe 7, P1	3 meters
U22	0B	U37	Storage server	PCIe 7, P1	3 meters
U22	1A	U35	Storage server	PCIe 7, P1	3 meters
U22	1B	U33	Storage server	PCIe 7, P1	3 meters
U22	2A	U31	Storage server	PCIe 7, P1	3 meters
U22	2B	U29	Storage server	PCIe 7, P1	3 meters
U22	3B	U27	Storage server	PCIe 7, P1	3 meters
U22	4A	U25	Storage server	PCIe 7, P1	2 meters
U22	5A	U23	Storage server	PCIe 7, P1	2 meters
U22	13A	U18	Storage server	PCIe 7, P2	2 meters
U22	13B	U17	Compute server	PCIe 7, P2	2 meters
U22	14B	U16	Compute server	PCIe 7, P2	2 meters
U22	14A	U14	Storage server	PCIe 7, P2	3 meters
U22	15A	U12	Storage server	PCIe 7, P2	3 meters
U22	15B	U10	Storage server	PCIe 7, P2	3 meters
U22	16A	U8	Storage server	PCIe 7, P2	3 meters
U22	16B	U6	Storage server	PCIe 7, P2	3 meters

Table F-16 (Cont.) RDMA Network FabricSwitch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	Device	Port	Cable Length
U22	17A	U4	Storage server	PCIe 7, P2	3 meters
U22	17B	U2	Storage server	PCIe 7, P2	3 meters
U20	6B	U41	Storage server	PCIe 7, P2	3 meters
U20	0A	U39	Storage server	PCIe 7, P2	3 meters
U20	0B	U37	Storage server	PCIe 7, P2	3 meters
U20	1A	U35	Storage server	PCIe 7, P2	3 meters
U20	1B	U33	Storage server	PCIe 7, P2	3 meters
U20	2A	U31	Storage server	PCIe 7, P2	3 meters
U20	2B	U29	Storage server	PCIe 7, P2	3 meters
U20	3B	U27	Storage server	PCIe 7, P2	3 meters
U20	4A	U25	Storage server	PCIe 7, P2	2 meters
U20	5A	U23	Storage server	PCIe 7, P2	2 meters
U20	13A	U18	Storage server	PCIe 7, P1	2 meters
U20	13B	U17	Compute server	PCIe 2, P1	2 meters
U20	14B	U16	Compute server	PCIe 2, P1	2 meters
U20	14A	U14	Storage server	PCIe 7, P1	2 meters
U20	15A	U12	Storage server	PCIe 7, P1	3 meters
U20	15B	U10	Storage server	PCIe 7, P1	3 meters
U20	16A	U8	Storage server	PCIe 7, P1	3 meters
U20	16B	U6	Storage server	PCIe 7, P1	3 meters
U20	17A	U4	Storage server	PCIe 7, P1	3 meters
U20	17B	U2	Storage server	PCIe 7, P1	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

G

In-Rack Cabling Tables for Recovery Appliance X4

This appendix identifies the location of the various components in a Recovery Appliance X4 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

This appendix contains the following sections:

The cables used in Recovery Appliance are color-coded as follows:

- **Black**: InfiniBand cables, or AC power jumper cables
- **Blue**: Gigabit Ethernet cables
- **Red**: Integrated Lights Out Manager (ILOM) cables

Recovery Appliance X4 Rack Layout

Figure G-1 shows the front and rear views of a Recovery Appliance X4 full rack. The front is shown on the left, and the rear is shown on the right.

> **Note:**
>
> The 1U empty space at the bottom of the rack is reserved for an optional InfiniBand spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure G-1 Rack Layout of the Recovery Appliance X4 Full Rack

Rear View

U#	PDU A	Zero Data Loss Recovery Appliance	PDU B
42	INP2	1U Solid Filler	INP0
41		1U Solid Filler	
40		Storage Server	
39			
38		Storage Server	
37			
36		Storage Server	
35			
34		Storage Server	
33			
32		Storage Server	
31			
30		Storage Server	
29			
28	INP1	Storage Server	INP1
27			
26		1U Solid Filler	
25		1U Solid Filler	
24		1U Solid Filler	
23		1U Solid Filler	
22		Infiniband Gateway Switch	
21		Ethernet Switch	
20		Infiniband Gateway Switch	
19		1U Solid Filler	
18		1U Solid Filler	
17		Compute Server	
16		Compute Server	
15		Storage Server	
14			
13	INP0	Storage Server	INP2
12			
11		Storage Server	
10			
9		Storage Server	
8			
7		Storage Server	
6			
5		Storage Server	
4			
3		Storage Server	
2			
1		1U Solid Filler	

Front View

U#	PDU B	Zero Data Loss Recovery Appliance	PDU A
42	INP0	1U Solid Filler	INP2
41		1U Solid Filler	
40		Storage Server	
39			
38		Storage Server	
37			
36		Storage Server	
35			
34		Storage Server	
33			
32		Storage Server	
31			
30		Storage Server	
29			
28	INP1	Storage Server	INP1
27			
26		1U Solid Filler	
25		1U Solid Filler	
24		1U Solid Filler	
23		1U Solid Filler	
22		1U Vented Filler	
21		1U Vented Filler	
20		1U Vented Filler	
19		1U Solid Filler	
18		1U Solid Filler	
17		Compute Server	
16		Compute Server	
15		Storage Server	
14			
13	INP2	Storage Server	INP0
12			
11		Storage Server	
10			
9		Storage Server	
8			
7		Storage Server	
6			
5		Storage Server	
4			
3		Storage Server	
2			
1		1U Solid Filler	

Figure G-2 shows the front and rear views of a Recovery Appliance X4 minimum configuration.

Figure G-2 Rack Layout of the Recovery Appliance X4 Minimum Configuration

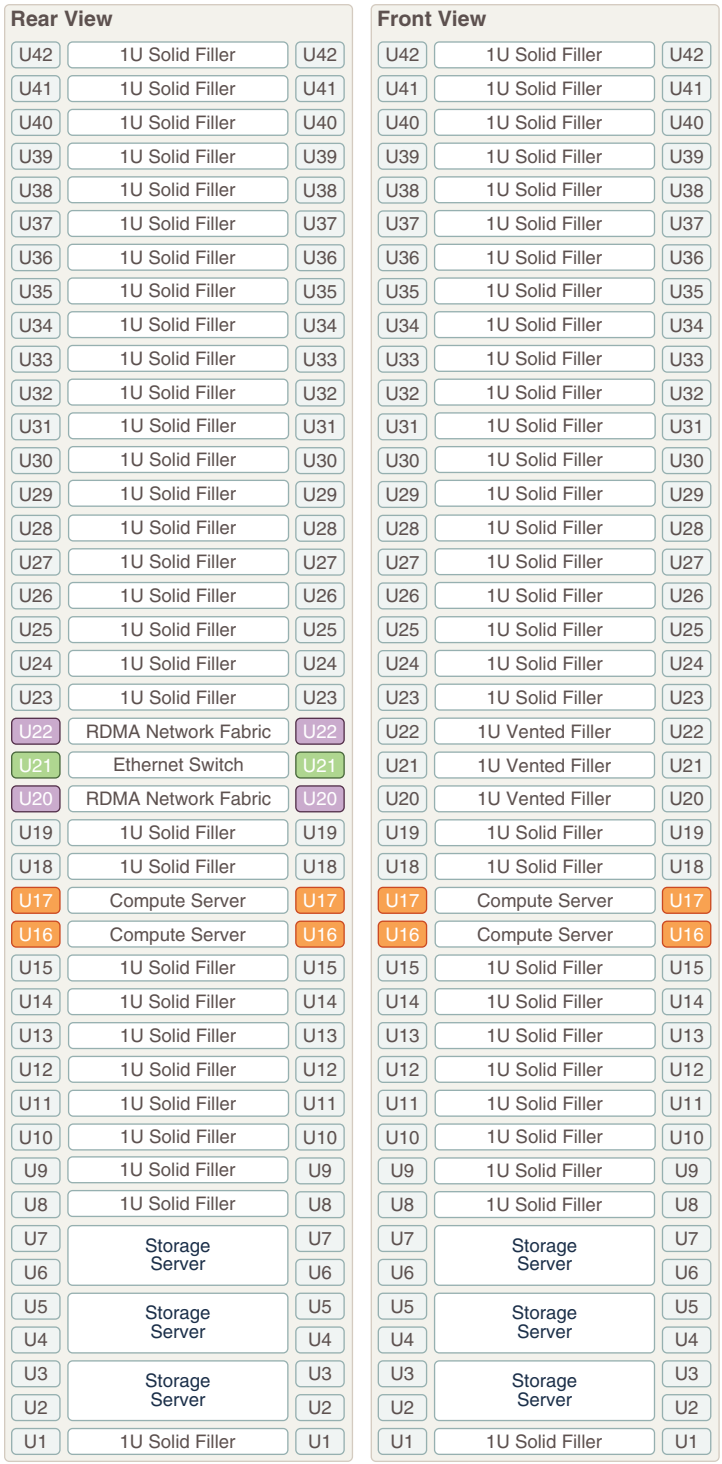

Oracle Integrated Lights Out Manager Cabling Connections

Table G-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X4 full rack. The port numbers are the same in racks with fewer storage

servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table G-1 ILOM Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

1 U*n* is the unit location in the rack, where *n* is the number.

Administrative Gigabit Ethernet Cabling Connections

Table G-2 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X4 full rack. The port numbers are the same in racks with fewer storage servers. The Ethernet switch is located in rack unit 21. Ethernet cables are blue.

Table G-2 Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U39	Storage server	1
U37	Storage server	3
U35	Storage server	5
U33	Storage server	7
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13

Table G-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit[1]	Component	To Ethernet Port
U22	InfiniBand leaf switch	45
U20	InfiniBand leaf switch	46
U17	Compute server	27
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

1 U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections

Table G-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X4 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table G-3 Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-0	G0-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G2-0	2 meters
U22	G3-1	G2-5	2 meters
U21	G3-0	G2-6	2 meters
U20	G2-4	G3-2	2 meters
U17	G2-1	G3-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters

Table G-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections

Table G-4 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X4 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table G-4 Three-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit[1]	PDU-A	PDU-B	Cable Length
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G0-0	2 meters
U22	G3-1	G0-5	2 meters
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U17	G2-1	G5-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G5-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters

1 U*n* is the unit location in the rack, where *n* is the number.

InfiniBand Switch-to-Switch Cabling Connections

Table G-5 lists the ports and cable lengths for the InfiniBand connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) InfiniBand cables.

Table G-5 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch U20 Port[1]	To InfiniBand Switch U22 Port[1]	Cable Length
9B	9A	1 meter
10B	10A	1 meter
11B	11A	1 meter
8A	8A	1 meter
9A	9B	1 meter
10A	10B	1 meter
11A	11B	1 meter

1 U*n* is the unit location in the rack, where *n* is the number.

InfiniBand Switch-to-Server Cabling Connections

Table G-6 lists the location, ports, and cables for the InfiniBand connections from the gateway leaf switches to the servers in a Recovery Appliance X4 full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) InfiniBand cables.

Table G-6 InfiniBand Switch-to-Server Cabling for the Recovery Appliance X4 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	To Rack Unit	Port	Cable Length
U22	0A	U39	Storage server	PCIe 3, P1	3 meters
U22	0B	U37	Storage server	PCIe 3, P1	3 meters
U22	1A	U35	Storage server	PCIe 3, P1	3 meters
U22	1B	U33	Storage server	PCIe 3, P1	3 meters
U22	2A	U31	Storage server	PCIe 3, P1	3 meters
U22	2B	U29	Storage server	PCIe 3, P1	3 meters
U22	3B	U27	Storage server	PCIe 3, P1	3 meters
U22	13B	U17	Compute server	PCIe 2, P2	2 meters
U22	14B	U16	Compute server	PCIe 2, P2	2 meters
U22	14A	U14	Storage server	PCIe 3, P2	3 meters

Table G-6 (Cont.) InfiniBand Switch-to-Server Cabling for the Recovery Appliance X4 Full Racks

From InfiniBand Switch Rack Unit[1]	Port	To Rack Unit	To Rack Unit	Port	Cable Length
U22	15A	U12	Storage server	PCIe 3, P2	3 meters
U22	15B	U10	Storage server	PCIe 3, P2	3 meters
U22	16A	U8	Storage server	PCIe 3, P2	3 meters
U22	16B	U6	Storage server	PCIe 3, P2	3 meters
U22	17A	U4	Storage server	PCIe 3, P2	3 meters
U22	17B	U2	Storage server	PCIe 3, P2	3 meters
U20	0A	U39	Storage server	PCIe 3, P2	3 meters
U20	0B	U37	Storage server	PCIe 3, P2	3 meters
U20	1A	U35	Storage server	PCIe 3, P2	3 meters
U20	1B	U33	Storage server	PCIe 3, P2	3 meters
U20	2A	U31	Storage server	PCIe 3, P2	3 meters
U20	2B	U29	Storage server	PCIe 3, P2	3 meters
U20	3B	U27	Storage server	PCIe 3, P2	3 meters
U20	13B	U17	Compute server	PCIe 2, P1	2 meters
U20	14B	U16	Compute server	PCIe 2, P1	2 meters
U20	14A	U14	Storage server	PCIe 3, P1	2 meters
U20	15A	U12	Storage server	PCIe 3, P1	3 meters
U20	15B	U10	Storage server	PCIe 3, P1	3 meters
U20	16A	U8	Storage server	PCIe 3, P1	3 meters
U20	16B	U6	Storage server	PCIe 3, P1	3 meters
U20	17A	U4	Storage server	PCIe 3, P1	3 meters
U20	17B	U2	Storage server	PCIe 3, P1	3 meters

1 U*n* is the unit location in the rack, where *n* is the number.

Index

A

B

C

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W